arXiv:1606.00136v1 [stat.ML] 1 Jun 2016

Efficiently Bounding Optimal Solutions after Small Data Modification in Large-Scale Empirical Risk Minimization

Hiroyuki Hanada
Nagoya Institute of Technology
Nagoya, Aichi, Japan
hanada.hiroyuki@nitech.ac.jp

Atsushi Shibagaki
Nagoya Institute of Technology
Nagoya, Aichi, Japan
shibagaki.a.mllab.nit@gmail.com

Jun Sakuma
University of Tsukuba
Tsukuba, Ibaraki, Japan
jun@cs.tsukuba.ac.jp

Ichiro Takeuchi
Nagoya Institute of Technology
Nagoya, Aichi, Japan
takeuchi.ichiro@nitech.ac.jp

June 1, 2016

Abstract

We study large-scale classification problems in changing environments where a small part of the dataset is modified, and the effect of the data modification must be quickly incorporated into the classifier. When the entire dataset is large, even if the amount of the data modification is fairly small, the computational cost of re-training the classifier would be prohibitively large. In this paper, we propose a novel method for efficiently incorporating such a data modification effect into the classifier without actually re-training it. The proposed method provides bounds on the unknown optimal classifier with the cost only proportional to the size of the data modification. We demonstrate through numerical experiments that the proposed method provides sufficiently tight bounds with negligible computational costs, especially when a small part of the dataset is modified in a large-scale classification problem.

1 Introduction

In this paper we study a problem of training a classifier such as the support vector machine (SVM) with a large-scale dataset. When a trained classifier is used in a changing environment where a part of the dataset is modified, we need to accordingly update the classifier for incorporating the effect of the data modification.

However, when the entire dataset is large, even if the amount of the data modification is fairly small, the computational cost of re-training the classifier would be prohibitively large. In such a situation of modifying only a tiny part of the dataset, we would expect that the optimal classifier trained with the modified dataset (called *new* classifier) would not be so different from the original classifier which was trained with the original dataset (called *old* classifier). Thus we might feel that spending a great amount of computational cost for re-training the classifier is not well-worthy effort, and feel like to use the old classifier without caring the effect of the data modification.

In this paper, we propose a novel efficient method for properly incorporating the effect of the data modification without actually re-training the classifier. The proposed method provides bounds on the new classifier with the cost only proportional to the size of the data modification. This computational advantage is beneficial especially when the size of the data modification is much smaller than the size of the entire dataset. The bounds obtained by the proposed method can be used for various tasks in classification problems. The bounds can be also used for computing the upper bound of the difference of the old and the new classifiers, which would be helpful for deciding when we should actually re-train the classifier.

Concretely, consider a linear binary classification problem with n instances and d features, and denote the data matrix as $X \in \mathbb{R}^{n \times d}$. Figure 1 illustrates the three scenarios considered in this paper. Denoting the set of modified elements in X as $\mathcal{M}$, we are particularly interested in the situation where only a small part of X is modified, i.e., $|\mathcal{M}|$ is much smaller than nd. Let $\hat{\boldsymbol{w}}^* \in \mathbb{R}^d$ and $\tilde{\boldsymbol{w}}^* \in \mathbb{R}^d$ be the linear model parameters of the old and new classifiers, respectively. The proposed method provides a lower bound $L(\tilde{w}_j)$ and an upper bound $U(\tilde{w}_j)$ such that

$$L(\tilde{w}_j^*) \leq \tilde{w}_j^* \leq U(\tilde{w}_j^*) \text{ for each } j = 1, \ldots, d \tag{1}$$

with the cost only proportional to the number of the modified elements $|\mathcal{M}|$.

Another novel contribution in this paper is to present methods for improving the bounds in (1) by what we call the *partial optimization*. When the number of features d is large, optimizing all the d coefficients $\{w_j\}_{j=1}^d$ would be very costly. For example, when $\mathcal{M} = \{\{(i, j')\}_{i=1}^n\}$, i.e., the j'^{th} column of X is modified (scenario (c) in Figure 1), we would conjecture that j'^{th} parameter would change a lot while other parameters would not change so much (unless the correlation of the feature with the j^{th} feature is very large). Thus we propose a partial optimization approach in which we only optimize a part of the parameters that are expected to change largely. A nice thing about the partial optimization approach is that the bounds in (1) can be improved with less computational cost than optimizing all the parameters. In the above example, we would optimize only j'^{th} parameter for computing tighter bounds in the form of (1). We also consider dual problems of linear classifier training where we have n dual parameters $\{\alpha_i\}_{i=1}^n$. In the case where i'^{th} row of the data matrix X is modified (scenario (b) in Figure 1), we can conduct partial optimization where only $\alpha_{i'}$ is optimized while other dual parameters are fixed. Although partial optimization-like approaches

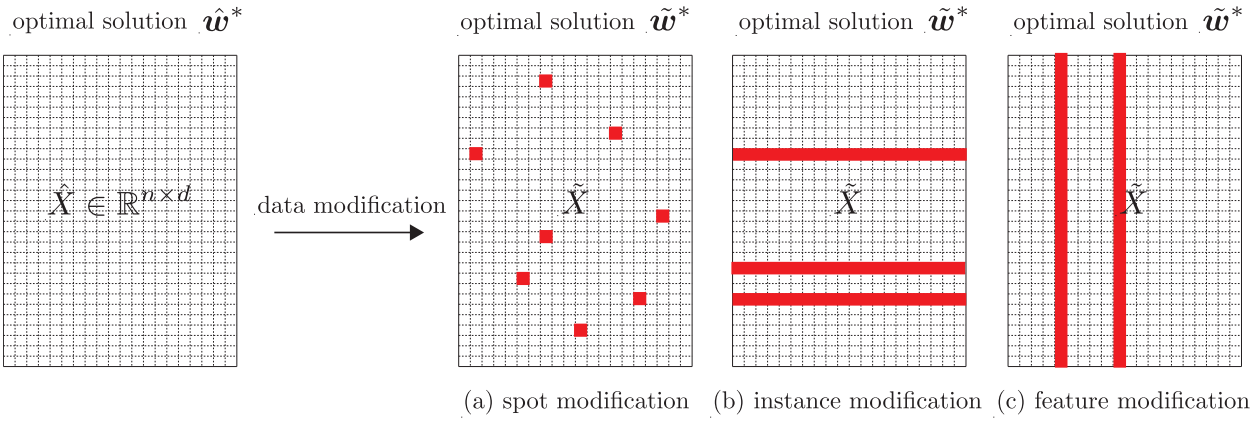

Figure 1: Schematic illustration of the three scenarios considered in this paper. When a small part of the data matrix X is modified, we need to quickly incorporate the effect of the data modification into the classifier. Our proposed method can compute a lower and an upper bound of the new optimal solution parameter $\tilde{w}_j^*$ with the cost only proportional to the number of modified elements.

might have been used in practice as a simple heuristic, to the best of our knowledge, there are no previous studies which points out that partial optimization can be useful for theoretically-sound non-heuristic decision making.

1.1 Related Works

In many practical machine learning tasks, we often need to solve multiple related optimization problems. In such a case, *warm-start* approach is sometimes useful where the optimal solution of a related problem is used as an initial starting point of the optimization problem [1,2]. In the data modification scenarios we consider here, we can use the old parameters $\hat{\boldsymbol{w}}^*$ as an initial starting point of the optimization problem for obtaining the new parameters $\tilde{\boldsymbol{w}}^*$. Unfortunately, however, even when we use warm-start, the computational cost of re-training an SVM-like classifier is at least $\mathcal{O}(nd)$ because we need to go through the entire data matrix X at least once in the optimization algorithm.

In optimization lieterature, so-called *sensitivity analysis* has been studied for the purpose of investigating the effect of *problem parameters* on the optimal solution, where problem parameters refer to the parameters that define the optimization problem. In the data modification scenarios we consider here, each element of the data matrix X can be considered as a problem parameter because it defines the optimization problem. Therefore, sensitivity analysis can be used for knowing how the classifier would change by the data modification in the vicinity of the old solution $\hat{\boldsymbol{w}}$. Unfortunately, however, sensitivity analysis cannot guarantee anything about the yet-to-be trained new solution $\tilde{\boldsymbol{w}}$.

The idea of bounding the solution of an optimization problem is inspired from the recent series of works on *safe screening* [3–11]. Safe screening is used for identifying non-active features or/and instances in sparse modeling. For example, in Lasso problem, activeness of each feature (whether the feature has nonzero

coefficient or not) is characterized by an inequality constraint on the dual variables. A key idea in safe screening is that a part of non-active features can be identified without actually solving the optimization problem if one can compute a bound of the dual optimal solution. In deriving the bounds in the form of (1), we use some techniques developed in safe screening literature.

In the context of incremental learning, a closely related approach was recently proposed in [12], in which the authors considered the case where a small number of instances are incrementally added or/and deleted from X (similar to scenario (b) in Figure1), and proposed an algorithm for computing the bounds of the new solution with the cost depending on $\Delta n \times d$, where Δn is the number of added and deleted instances. Although their bound formula is totally different from ours, in the scenario on instance modification, the time complexities for computing their bounds and our bounds are same. However, the method in [12] cannot be used for other scenarios such as scenario (a) or (c) in Figure1. Furthermore, as we clarify in §3, our method can be applied to wider class of learning problems than the one considered in [12]. The work in [13] also studied similar problems in distributed stream learning context. In this context, the authors in [13] studied how to efficiently compute an upper bound of the difference between the old and the current solutions.

To the best of our knowledge, there are no other existing studies that study the above problem setup in this general form. In the current big data era, it is quite important to be able to quickly incorporate the effect of a data modification in a huge database without going through the entire database.

2 Problem Setup and Background

2.1 Notations

For any natural number n, we define $[n] := \{1, \dots, n\}$. We write the subdifferential operator as ∂. For a function f, its domain is denoted as $\mathrm{dom} f$, while its conjugate function is written as f^*. The L_2 norm of a vector $\boldsymbol{v} \in \mathbb{R}^m$ is written as $\|\boldsymbol{v}\|_2 := (\sum_{k\in[m]} |v_k|^2)^{1/2}$.

2.2 Convex Regularized Learning Problems

In this paper, we study binary classification problems with n training instances and d features. The entire dataset is denoted as $(X, \boldsymbol{y})$, where $X \in \mathbb{R}^{n\times d}$ is the input matrix and $\boldsymbol{y} \in \{\pm 1\}^n$ is the label vector. The i^{th} row vector, the j^{th} column vector, and $(i,j)^{\text{th}}$ element of X are respectively written as $\boldsymbol{x}_{i\cdot} \in \mathbb{R}^d$, $\boldsymbol{x}_{\cdot j} \in \mathbb{R}^n$, and $x_{ij} \in \mathbb{R}$. Furthermore, we define $Z := \mathrm{diag}(\boldsymbol{y})X \in \mathbb{R}^{n\times d}$ for notational simplicity, and $\boldsymbol{z}_{i\cdot} \in \mathbb{R}^d$, $\boldsymbol{z}_{\cdot j} \in \mathbb{R}^n$, $z_{ij} \in \mathbb{R}$ are also defined similarly. We consider the following class of convex regularized empirical risk minimization problems:

$$\boldsymbol{w}^* = \arg\min_{\boldsymbol{w}\in\mathbb{R}^d} \mathcal{P}(\boldsymbol{w}) := \frac{1}{n}\sum_{i\in[n]} \phi(\boldsymbol{z}_{i\cdot}^\top \boldsymbol{w}) + \psi(\boldsymbol{w}), \tag{2}$$

where $\phi : \mathbb{R} \to \mathbb{R}_+$ is a convex loss function and $\psi : \mathbb{R}^d \to \mathbb{R}_+$ is a convex penalty function. Using Fenchel's duality theorem (see, e.g., Corollary 31.2.1 in [14]), the dual problem of (2) is written as

$$\boldsymbol{\alpha}^* := \arg\max_{\boldsymbol{\alpha} \in \mathbb{R}^n} \mathcal{D}(\boldsymbol{\alpha}) := -\frac{1}{n} \sum_{i \in [n]} \phi^*(-\alpha_i) - \psi^* \left(\frac{1}{n} Z^\top \boldsymbol{\alpha} \right), \tag{3}$$

where ϕ^* and ψ^* are convex conjugates of ϕ and ψ, respectively.

Example (Smoothed-hinge SVM) As a working example, we study smoothed hinge SVM which is obtained by setting

$$\phi(r) := \begin{cases} 0 & (r > 1), \\ 1 - r - \frac{\gamma}{2} & (r < 1 - \gamma), \\ \frac{1}{2\gamma}(1-r)^2 & (\text{otherwise}), \end{cases} \qquad \psi(\boldsymbol{w}) := \frac{\lambda}{2} \|\boldsymbol{w}\|_2^2, \tag{4}$$

where $\gamma > 0$ and $\lambda > 0$ are tuning parameters. Their convex conjugate functions are respectively written as

$$\phi^*(r) = \begin{cases} r + \frac{\gamma}{2} r^2 & (-1 \leq r \leq 0), \\ +\infty & (\text{otherwise}), \end{cases} \qquad \psi^*(\boldsymbol{w}) = \frac{1}{2\lambda} \|\boldsymbol{w}\|_2^2. \tag{5}$$

2.3 Small Data Modification

In this study, we consider a situation that a small portion of the input matrix X is modified. The set of modified elements in X is denoted as $\mathcal{M} \subseteq \{[n] \times [d]\}$ and its size as $|\mathcal{M}|$. Furthermore, we define $\mathcal{M}_i := \{i \in [n] \mid \exists j \in [d] \text{ s.t. } (i,j) \in \mathcal{M}\}$ and $\mathcal{M}_j := \{j \in [d] \mid \exists i \in [n] \text{ s.t. } (i,j) \in \mathcal{M}\}$. We call the problems before and after the data modification as *old* problem and *new* problem, respectively. We use hat $\hat{}$ notation for the old problem and tilde $\tilde{}$ notation for the new problem; e.g., $\hat{X}$, $\hat{\boldsymbol{x}}_{i\cdot}$, $\hat{\boldsymbol{x}}_{\cdot j}$, $\hat{x}_{ij}$, $\hat{Z}$, $\hat{\boldsymbol{z}}_{i\cdot}$, $\hat{\boldsymbol{z}}_{\cdot j}$, $\hat{z}_{ij}$, etc. represent the old data before the modification, while $\tilde{X}$, $\tilde{\boldsymbol{x}}_{i\cdot}$, $\tilde{\boldsymbol{x}}_{\cdot j}$, $\tilde{x}_{ij}$, $\tilde{Z}$, $\tilde{\boldsymbol{z}}_{i\cdot}$, $\tilde{\boldsymbol{z}}_{\cdot j}$, $\tilde{z}_{ij}$, etc. represent the new data after the modification. Furthermore, the primal and the dual optimal solutions for the old and the new problems are respectively denoted as $(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)$ and $(\tilde{\boldsymbol{w}}^*, \tilde{\boldsymbol{\alpha}}^*)$.

Our main contribution in this paper is to present methods for computing bounds of each primal variable $\tilde{w}_j^*, j \in [d]$ and dual variable $\tilde{\alpha}_i^*, i \in [n]$ with the computational cost depending only on the number of the modified elements $|\mathcal{M}|$. A lower bound and an upper bound of a scalar quantity are respectively represented by using a capital letter L and a capital letter U such as

$$L(\tilde{w}_j^*) \leq \tilde{w}_j^* \leq U(\tilde{w}_j^*), j \in [d], \qquad L(\tilde{\alpha}_i^*) \leq \tilde{\alpha}_i^* \leq U(\tilde{\alpha}_i^*), i \in [n].$$

2.4 Decision Making by New Solution's Bounds

Using these bounds, we can perform several decision making tasks on the new solution $(\tilde{\boldsymbol{w}}^*, \tilde{\boldsymbol{\alpha}}^*)$.

First, consider a situation that we want to classify a test instance $\boldsymbol{x}' \in \mathbb{R}^d$ based on the new classifier. In this task, we compute a lower bound and an upper bound of the classification score $f(\boldsymbol{x}') = \boldsymbol{x}'^\top \tilde{\boldsymbol{w}}^*$. Using the lower and the upper bounds, the class label y' of the test instance $\boldsymbol{x}'$ can be simply determined as

$$y' = \mathrm{sign}(\boldsymbol{x}'^\top \tilde{\boldsymbol{w}}^*) = \begin{cases} -1 & (U(\boldsymbol{x}'^\top \tilde{\boldsymbol{w}}^*) < 0), \\ +1 & (L(\boldsymbol{x}'^\top \tilde{\boldsymbol{w}}^*) \geq 0), \\ \text{unknown} & (\text{otherwise}). \end{cases} \tag{6}$$

Equation (6) suggests that, even if the new optimal solution $\tilde{\boldsymbol{w}}^*$ is not available, if we have its sufficiently tight bounds, the classification task can be completed for some instances[1]. Similar discussion has been done in [12] and [15].

Next, consider a situation that the data matrix X is constantly changing. In such a situation, an important decision making task is to determine when we should actually re-train the classifier. Suppose that we have a tolerance threshold $\theta > 0$ for how much the classifier can be different from the optimal one. If we quantify the difference by L_2 norm $\|\hat{\boldsymbol{w}}^* - \tilde{\boldsymbol{w}}^*\|_2$, we can re-train the classifier only when

$$U(\|\hat{\boldsymbol{w}}^* - \tilde{\boldsymbol{w}}^*\|_2) \geq \theta. \tag{7}$$

Equation (7) indicates that, even if the optimal solution $\tilde{\boldsymbol{w}}^*$ is not available, we can make sure that the difference of the currently available solution $\hat{\boldsymbol{w}}^*$ from the unknown optimal solution $\tilde{\boldsymbol{w}}^*$ is no greater than our tolerance threshold[2]. Similar discussion has been done in [12] and [13].

Finally, let us discuss the connection with safe screening [3–11]. In the example of smoothed hinge SVM (see (4) and (5)), the optimality condition indicates that $1 < \tilde{\boldsymbol{z}}_i^\top \tilde{\boldsymbol{w}}^* \;\Rightarrow\; \tilde{\alpha}_i^* = 0$. This condition indicates that $L(\tilde{\boldsymbol{z}}_i^\top \tilde{\boldsymbol{w}}^*) > 1 \;\Rightarrow\; \tilde{\alpha}_i^* = 0$, which means that the i^{th} instance would be non-support vector, and we can exclude the instance before solving the new optimization problem. This approach is referred to as *safe sample screening* [16–18].

[1] Using $L(\tilde{w}_j)$ and $U(\tilde{w}_j)$ for $j \in [d]$, we can compute a lower bound and an upper bound of $\boldsymbol{x}'^\top \tilde{\boldsymbol{w}}$ as

$$L(\boldsymbol{x}'^\top \tilde{\boldsymbol{w}}^*) = \sum_{j|x'_{ij} \geq 0} x'_{ij} L(\tilde{w}_j^*) + \sum_{j|x'_{ij} < 0} x'_{ij} U(\tilde{w}_j^*), U(\boldsymbol{x}'^\top \tilde{\boldsymbol{w}}^*) = \sum_{j|x'_{ij} \geq 0} x'_{ij} U(\tilde{w}_j^*) + \sum_{j|x'_{ij} < 0} x'_{ij} L(\tilde{w}_j^*).$$

[2] Using $L(\tilde{w}_j)$ and $U(\tilde{w}_j)$ for $j \in [d]$, we can compute an upper bound of $\|\hat{\boldsymbol{w}} - \tilde{\boldsymbol{w}}\|_2$ as

$$U(\|\hat{\boldsymbol{w}} - \tilde{\boldsymbol{w}}\|_2) = \sqrt{\sum_{j \in [d]} \max\{\hat{w}_j^* - L(\tilde{w}_j^*), U(\tilde{w}_j^*) - \hat{w}_j^*\}^2}.$$

3 Efficient Bound Computation after Small Data Change

We present our main results in this section. Our goal is to efficiently compute lower and upper bounds of the new primal-dual optimal solutions $\tilde{\boldsymbol{w}}^*$ and $\tilde{\boldsymbol{\alpha}}^*$ by using the old primal-dual optimal solutions $(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)$. If the loss function ϕ or/and the penalty function ψ have certain properties, we show that lower and upper bounds of $\tilde{w}_j^*, j \in [d]$ and $\tilde{\alpha}_i^*, i \in [n]$ can be computed with time complexity $\mathcal{O}(|\mathcal{M}|)$ where $|\mathcal{M}|$ is the number of modified elements. When the number of the modified elements is much smaller than the number of the entire elements nd, it is quite beneficial to be able to compute the bounds of the new optimal solution in such an efficient way.

Before presenting the main theorem, we summarize the properties of the loss function and the penalty function:

Property A (for loss function): The loss function ϕ is γ^{-1}-smooth, i.e., ϕ is differentiable and its gradient $\partial\phi$ satisfies

$$|\partial\phi(a) - \partial\phi(b)| \leq \gamma^{-1}|a - b| \quad \forall a, b \in \mathbb{R}.$$

Property B (for penalty function): The penalty function ψ is λ-strongly convex, i.e., ψ is sub-differentiable and its sub-gradient $\partial\psi$ satisfies

$$\psi(\boldsymbol{a}) \geq \psi(\boldsymbol{b}) + \partial\psi(\boldsymbol{b})^\top(\boldsymbol{a} - \boldsymbol{b}) + \frac{\lambda}{2}\|\boldsymbol{a} - \boldsymbol{b}\|_2^2 \quad \forall \boldsymbol{a}, \boldsymbol{b} \in \mathbb{R}^d.$$

At the end of this section, we will present examples of the loss functions and the penalty functions that possess these properties. The following theorem tells that, when the loss function and/or the penalty function have these properties, lower and upper bounds of the new solutions $\tilde{w}_j^*, j \in [d]$, and $\tilde{\alpha}_i^*, i \in [n]$ can be efficiently computed.

Theorem 1. *Assume that the following quantities*

$$\hat{\boldsymbol{z}}_{i\cdot}^\top \hat{\boldsymbol{w}},\ \|\hat{\boldsymbol{x}}_{i\cdot}\|_2,\ \hat{\boldsymbol{z}}_{\cdot j}^\top \hat{\boldsymbol{\alpha}},\ \|\hat{\boldsymbol{x}}_{\cdot j}\|_2,\ \inf_{\boldsymbol{\alpha} \in \mathrm{dom} D(\boldsymbol{\alpha})} n^{-1}\hat{\boldsymbol{z}}_{\cdot j}^\top \boldsymbol{\alpha},\ \sup_{\boldsymbol{\alpha} \in \mathrm{dom} D(\boldsymbol{\alpha})} n^{-1}\hat{\boldsymbol{z}}_{\cdot j}^\top \boldsymbol{\alpha}, \tag{8}$$

for all $i \in [n]$ and $j \in [d]$, are stored in the memory with the size $\mathcal{O}(n + d)$, and the penalty function ψ is decomposable *in the sense that there exists d convex functions $\psi_j : \mathbb{R} \to \mathbb{R}$, $j \in [d]$, such that*

$$\psi(\boldsymbol{w}) = \sum_{j \in [d]} \psi_j(w_j) \text{ for all } \boldsymbol{w} \in \mathbb{R}^d.$$

Then, in each of the following three cases (i), (ii), and (iii), the lower and the upper bounds of $\{\tilde{w}_j^*\}_{j \in [d]}$ and $\{\tilde{\alpha}_i^*\}_{i \in [n]}$ can be evaluated with time complexity $\mathcal{O}(|\mathcal{M}|)$.

Case (i) The loss function has **Property A***:*

$$\tilde{w}_j^* \geq L_D(\tilde{w}_j^*) := \inf \partial\psi_j^* \left(\max \left\{ \inf_{\boldsymbol{\alpha} \in \mathrm{dom}\mathcal{D}(\boldsymbol{\alpha})} n^{-1} \tilde{\boldsymbol{z}}_{\cdot j}^\top \boldsymbol{\alpha}, n^{-1} \tilde{\boldsymbol{z}}_{\cdot j}^\top \hat{\boldsymbol{\alpha}}^* - \|\tilde{\boldsymbol{x}}_{\cdot j}\|_2 \sqrt{2\gamma^{-1} \tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)} \right\} \right),$$

$$\tilde{w}_j^* \leq U_D(\tilde{w}_j^*) := \sup \partial\psi_j^* \left(\min \left\{ \sup_{\boldsymbol{\alpha} \in \mathrm{dom}\mathcal{D}(\boldsymbol{\alpha})} n^{-1} \tilde{\boldsymbol{z}}_{\cdot j}^\top \boldsymbol{\alpha}, n^{-1} \tilde{\boldsymbol{z}}_{\cdot j}^\top \hat{\boldsymbol{\alpha}}^* + \|\tilde{\boldsymbol{x}}_{\cdot j}\|_2 \sqrt{2\gamma^{-1} \tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)} \right\} \right),$$

$$\tilde{\alpha}_i^* \geq L_D(\tilde{\alpha}_i^*) := \min \left\{ \inf_{\boldsymbol{\alpha} \in \mathrm{dom}\mathcal{D}(\boldsymbol{\alpha})} \alpha_i, \hat{\alpha}_i - \sqrt{2n\gamma^{-1} \tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)} \right\},$$

$$\tilde{\alpha}_i^* \leq L_D(\tilde{\alpha}_i^*) := \max \left\{ \sup_{\boldsymbol{\alpha} \in \mathrm{dom}\mathcal{D}(\boldsymbol{\alpha})} \alpha_i, \hat{\alpha}_i + \sqrt{2n\gamma^{-1} \tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)} \right\},$$

where

$$\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*) := n^{-1} \sum_{i \in \mathcal{M}_i} \left(\phi(\tilde{\boldsymbol{z}}_{i\cdot}^\top \hat{\boldsymbol{w}}^*) - \phi(\hat{\boldsymbol{z}}_{i\cdot}^\top \hat{\boldsymbol{w}}^*) \right) + \sum_{j \in \mathcal{M}_j} \left(\psi_j^*(n^{-1} \tilde{\boldsymbol{z}}_{\cdot j}^\top \hat{\boldsymbol{\alpha}}^*) - \psi_j^*(n^{-1} \hat{\boldsymbol{z}}_{\cdot j}^\top \hat{\boldsymbol{\alpha}}^*) \right).$$

Case (ii) The penalty function has **Property B***:*

$$\tilde{w}_j^* \geq L_P(\tilde{w}_j^*) := \hat{w}_j - \sqrt{2\lambda^{-1} \tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)}, \quad \tilde{w}_j^* \leq U_P(\tilde{w}_j^*) \qquad := \hat{w}_j + \sqrt{2\lambda^{-1} \tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)},$$

$$\tilde{\alpha}_i^* \geq L_P(\tilde{\alpha}_i^*) := \inf \left(-\partial\phi \left(\tilde{\boldsymbol{z}}_{i\cdot}^\top \hat{\boldsymbol{w}}^* + \|\tilde{\boldsymbol{x}}_{\cdot i}\|_2 \sqrt{2\lambda^{-1} \tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)} \right) \right),$$

$$\tilde{\alpha}_i^* \leq U_P(\tilde{\alpha}_i^*) := \sup \left(-\partial\phi \left(\tilde{\boldsymbol{z}}_{i\cdot}^\top \hat{\boldsymbol{w}}^* - \|\tilde{\boldsymbol{x}}_{i\cdot}\|_2 \sqrt{2\lambda^{-1} \tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)} \right) \right).$$

Case (iii) The loss function has **Property A** *and the penalty function has* **Property B***:*

$$\tilde{w}_j^* \geq L_{PD}(\tilde{w}_j^*) := \max \left\{ L_P(\tilde{w}_j^*), L_D(\tilde{w}_j^*) \right\}, \quad \tilde{w}_j^* \leq U_{PD}(\tilde{w}_j^*) := \min \left\{ U_P(\tilde{w}_j^*), U_D(\tilde{w}_j^*) \right\},$$

$$\tilde{\alpha}_i^* \geq L_{PD}(\tilde{\alpha}_i^*) := \max \left\{ L_P(\tilde{\alpha}_i^*), L_D(\tilde{\alpha}_i^*) \right\}, \quad \tilde{\alpha}_i^* \leq U_{PD}(\tilde{\alpha}_i^*) := \min \left\{ U_P(\tilde{\alpha}_i^*), U_D(\tilde{\alpha}_i^*) \right\}.$$

The proof of Theorem 1 is presented in Appendix A. In the proof, we use a technique recently introduced in the context of safe screening based on the duality gap of an approximate solution [10, 11]. Note that the tightness of the bounds in the theorem depends on a quantity $\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)$, which indicates the duality gap of the new optimization problem evaluated at the old primal-dual solutions $\hat{\boldsymbol{w}}^*$ and $\hat{\boldsymbol{\alpha}}^*$. For proving the time complexity in the theorem, we exploit a simple fact that the duality gap $\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)$ can be computed with time complexity $\mathcal{O}(|\mathcal{M}|)$[3]. The assumption that all the quantities in (8) are available in the memory is

[3]Note that the duality gap of the old problem evaluated at $\hat{\boldsymbol{w}}^*$ and $\hat{\boldsymbol{\alpha}}^*$ is zero because of their optimality. If the duality gap of the old problem evaluated at $\hat{\boldsymbol{w}}^*$ and $\hat{\boldsymbol{\alpha}}^*$ is not exactly zero due to numerical issue, we can simply add the remaining gap to $\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)$.

reasonable because they can be computed when the old optimization problem was solved for obtaining $\hat{\boldsymbol{w}}^*$ and $\hat{\boldsymbol{\alpha}}^*$. The assumption on decomposability is satisfied in many penalty functions including L_q^q-norm with $q \geq 1$.

Many popular loss functions for binary classification problems such as squared hinge loss and logistic regression loss have **Property A**. Penalty function that has **Property B** includes L_2 penalty and elastic-net penalty, etc. Since we can efficiently compute bounds if either **Property A** or **Property B** is satisfied, we can, for example, compute bounds for standard vanilla-hinge SVM, L_1-penalized logistic regression, and many other popular classification problems. Similar bound computation methods can be derived also for a wide class of regularized convex regression problems.

The bounds in Theorem 1 for smoothed-hinge SVM are presented in Appendix B.

4 Partial Optimization for Obtaining Tighter Bounds

Let us consider how to handle a situation where the bounds in Theorem 1 are not sufficiently tight enough. If we do not mind spending computational cost depending on the entire data matrix size nd, the most naive solution would be to just re-train the classifier with the modified new dataset. However, when n and d are very large, it would be desirable to be able to obtain tighter bounds than those in Theorem 1 with a little more additional cost, but still not depending on nd.

Our idea is to simply note that the tightness of the bounds in Theorem 1 depends on the *closeness* of the old solution $(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)$ to the new solutions $(\tilde{\boldsymbol{w}}^*, \tilde{\boldsymbol{\alpha}}^*)$. It suggests that, if we could replace the old solution with another solution that is closer to the new solution, the bounds could be tightened. We propose a simple method for obtaining such a closer solution with the computational cost not depending on the entire data matrix size, and show that the method always improve the bounds in Theorem 1.

A primal variable w_j, $j \in [d]$, would be highly dependent on the j^{th} column of X. It suggests that, if many elements in the j^{th} column of X are modified, the j^{th} primal variable w_j would change a lot after the data modification, i.e., the difference between $\hat{w}_j^*$ and $\tilde{w}_j^*$ would be large. Similarly, a dual variable α_i, $i \in [n]$, would be highly affected by the i^{th} row of X. If many elements in the i^{th} row of X are modified, the difference between $\hat{\alpha}_i^*$ and $\tilde{\alpha}_i^*$ would be large.

These observations suggest that we might be able to obtain a solution closer to the new solution with reasonable computational cost by only optimizing a part of the variables that are expected to be highly affected by the data modification, while other variables are fixed. If we consider solving the primal optimization problem only w.r.t. a subset of the primal variables $w_j, j \in \mathcal{J} \subseteq [d]$, the computational cost depends only on $n|\mathcal{J}|$. Similarly, if we consider solving the dual optimization problem only w.r.t. a subset of the dual variables $\alpha_i, i \in \mathcal{I} \subseteq [n]$, the computational cost depends only on $d|\mathcal{I}|$. The following theorem tells that the bounds in Theorem 1 can be strictly tightened if these partial optimization problems have a strict progress in terms of their objective function values.

Theorem 2. *Assume that the penalty function is decomposable as defined in Theorem 1. Consider partially optimizing the primal problem* (2) *only w.r.t. a subset of the primal variables* $\{w_j\}_{j\in\mathcal{J}\subseteq[d]}$. *Let* $\check{\boldsymbol{w}}^*$ *be a* d*-dimensional vector whose elements corresponding to* $\mathcal{J}$ *are the solution of the partial primal problem, while the elements corresponding to* $[d]\setminus\mathcal{J}$ *are the old primal solutions* $\{\hat{w}_j^*\}_{j\in[d]\setminus\mathcal{J}}$. *Then, unless* $\check{\boldsymbol{w}}^*=\hat{\boldsymbol{w}}^*$, *all the bounds in Theorem 1 will be strictly tightened by replacing* $\hat{\boldsymbol{w}}^*$ *with* $\check{\boldsymbol{w}}^*$.

Similarly, consider partially optimizing the dual problem (3) *only w.r.t. a subset of the dual variables* $\{\alpha_i\}_{i\in\mathcal{I}\subseteq[n]}$. *Let* $\check{\boldsymbol{\alpha}}^*$ *be an* n*-dimensional vector whose elements corresponding to* $\mathcal{I}$ *are the solution of the partial dual problem, while the elements corresponding to* $[n]\setminus\mathcal{I}$ *are the old dual solutions* $\{\hat{\alpha}_i^*\}_{i\in[n]\setminus\mathcal{I}}$. *Then, unless* $\check{\boldsymbol{\alpha}}^*=\hat{\boldsymbol{\alpha}}^*$, *all the bounds in Theorem 1 will be strictly tightened by replacing* $\hat{\boldsymbol{\alpha}}^*$ *with* $\check{\boldsymbol{\alpha}}^*$.

The proof of Theorem 2 is presented in Appendix A.

For each of the three scenarios depicted in Figure 1, we consider specific partial optimization strategy. First, for spot modification scenario, we set $\mathcal{J}:=\mathcal{M}_j$ and $\mathcal{I}:=\mathcal{M}_i$, i.e., we optimized the primal (resp. dual) variables if there are at least one modification at the corresponding columns (resp. rows) of X. For instance modification scenario, we set $\mathcal{J}:=\emptyset$ and $\mathcal{I}:=\mathcal{M}_i$, i.e., we only optimized dual variables corresponding to the modified rows of X. Similarly for feature modification scenario, we set $\mathcal{J}:=\mathcal{M}_j$ and $\mathcal{I}:=\emptyset$, i.e., we only optimized primal variables corresponding to the modified columns of X.

5 Experiment

In this section, we empirically investigate the tightness of the bounds obtained by the approaches in §3 and §4.

5.1 Settings

We evaluated the tightness of the bounds in terms of the performances of the two tasks described in §2. The first task is test instance classification, where the predicted test label y' for a test instance $\boldsymbol{x}'\in\mathbb{R}^d$ can be determined by the bounds of the classification score $L(\boldsymbol{x}'^\top\tilde{\boldsymbol{w}}^*)$ and $U(\boldsymbol{x}'^\top\tilde{\boldsymbol{w}}^*)$ as in (6). To see the performances of this task, we report the rate of the test instances whose predicted class labels are determined. The second task is to evaluate how much the solution can change by the data modification. To see the performances of this task, we report the upper bound of the parameter change in terms of the L_2 norm.

We studied smoothed-hinge SVM in (4) and (5) on five benchmark datasets obtained from libsvm data repository [19] (Table 1), where 80% of the instances are used as training set and the rest are used as test set. We considered three scenarios in Figure 1. In spot modification scenario, we modified randomly chosen $|\mathcal{M}|$ elements in X for $|\mathcal{M}|\in\{1,100,10000\}$. In instance modification scenario, we randomly chose Δn rows of X for $\Delta n\in\{1,10,100\}$, and modified all non-zero values in them. In feature modification scenario, we randomly chose Δd columns of X for $\Delta d\in\{1,10,100\}$, and modified all non-zero values in them. To modify

Table 1: Data sets used for the experiment

- All are from LIBSVM data repository [19].
- 80% of instances are used for training while others are for validation.
- Every instance $\boldsymbol{x}$ is normalized to norm $\|\boldsymbol{x}\|_2 = 1$.

Name	n	d
`kdd-a`	8,407,752	20,216,830
`url`	2,396,130	3,231,961
`news20`	19,996	1,355,191
`real-sim`	72,309	20,958
`rcv1-train`	20,242	47,236

a value, we replaced it with a random value taken from the uniform distribution between the minimum and the maximum value of the corresponding feature in the data set. We set $\lambda = \{0.001, 0.01, 0.1, 1\}$. For each condition, we run the experiment with 10 different random seeds.

5.2 Results

Figures 2 to 5 show the results of the label determination rate, Figures 6 to 9 show the upper bound of the parameter change, and Table 2 show the ratio of the computation time of bound computation in §3 to re-training.

The red points in the figures show the performances of the bounds computed by our proposed method in §3. For `kdd-a` dataset, more than 99.9% of the test labels were determined by using the bounds, and the upper bounds of the parameter change were also highly stable. We conjecture that, since this dataset is huge ($n > 8$ million and $d > 20$ million), small amount of data modification did not change the solution much, and our bounds could nicely capture this phenomenon. For `url` dataset, the label determination rates were slightly worse especially when 100 rows were modified in instance modification scenario. For smaller datasets, we can confirm that the rates become worse. We conjecture that this performance deterioration might happen when many influential row vectors were modified and the change of the solution was relatively large. Table 2 compares the computational costs of our bound computation method with those of re-training (with warm-start). From the table, we observe that the cost of our bound computation method is almost negligible.

The blue points in the figures show the performances of the bounds lifted-up by partial optimization approach discussed in §4. The black lines connecting red and blue points indicate the correspondence in 10 random trials. As described in Theorem 2, the bounds in blue points are always tighter than those in red points (higher value in label determination rate and lower values in parameter change indicate tighter bounds). For `kdd-a` dataset, the performance of the red and blue points are almost same because the original bounds in red points were already tight enough. For other datasets, the partial optimization approach seemed to work well. Especially when the performances of the original bounds in red points were relatively poor (e.g., $\Delta n = 100$ case in instance modification scenario), the improvements by partial optimization were significant.

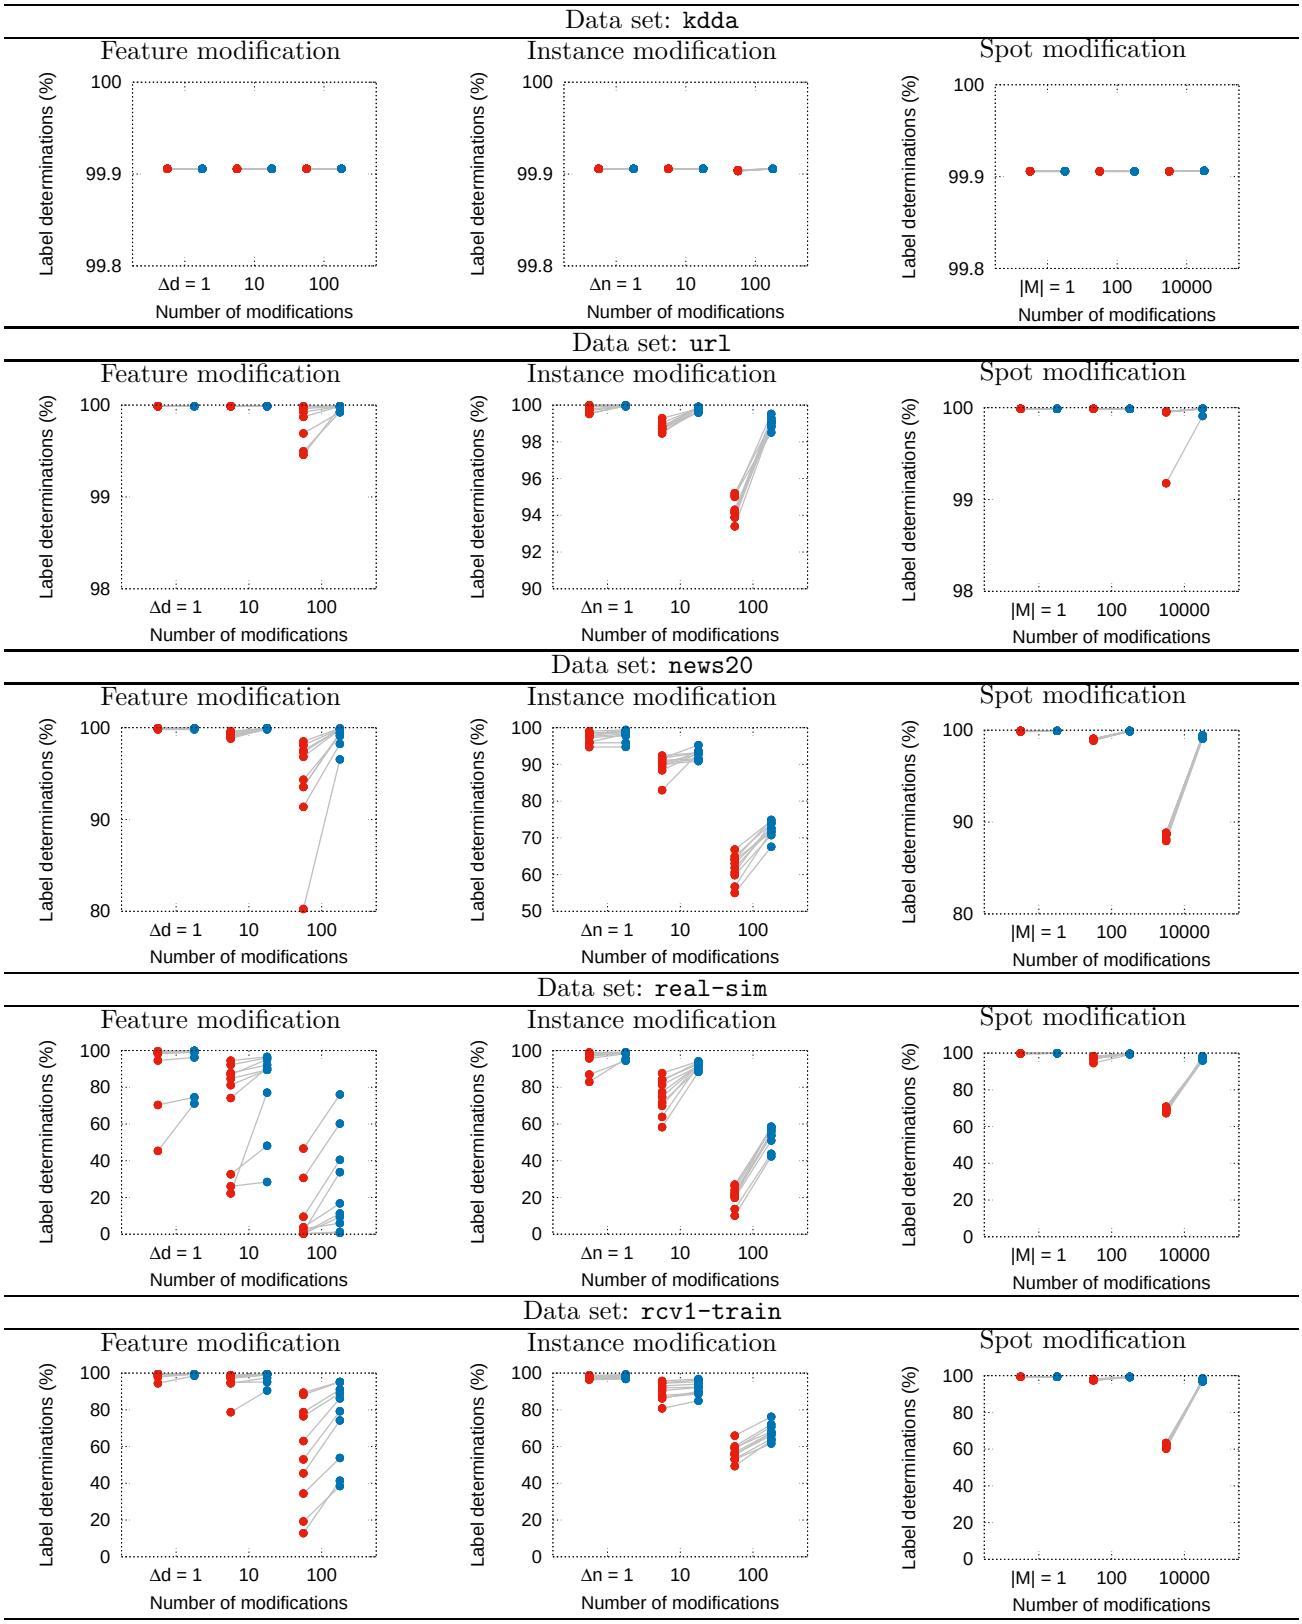

Figure 2: Label determination rate for $\lambda = 0.001$

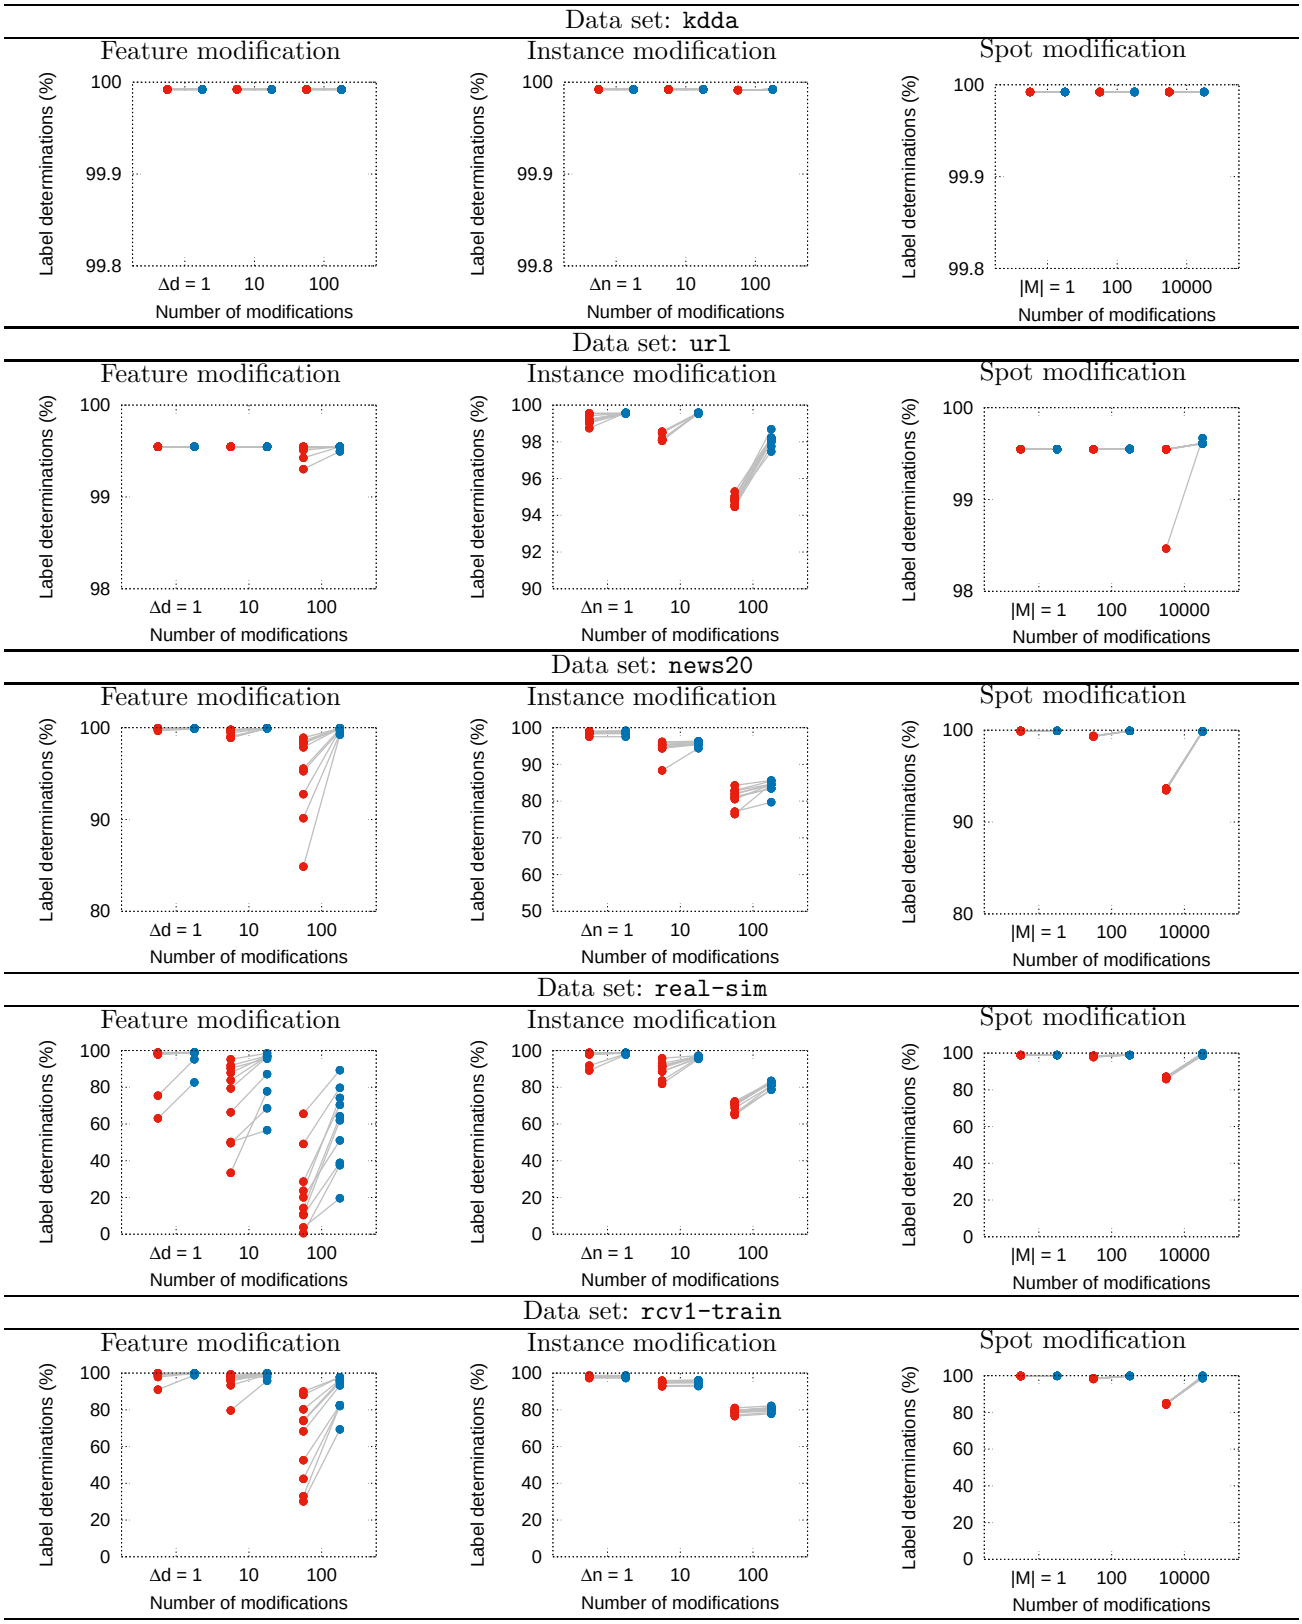

Figure 3: Label determination rate for $\lambda = 0.01$

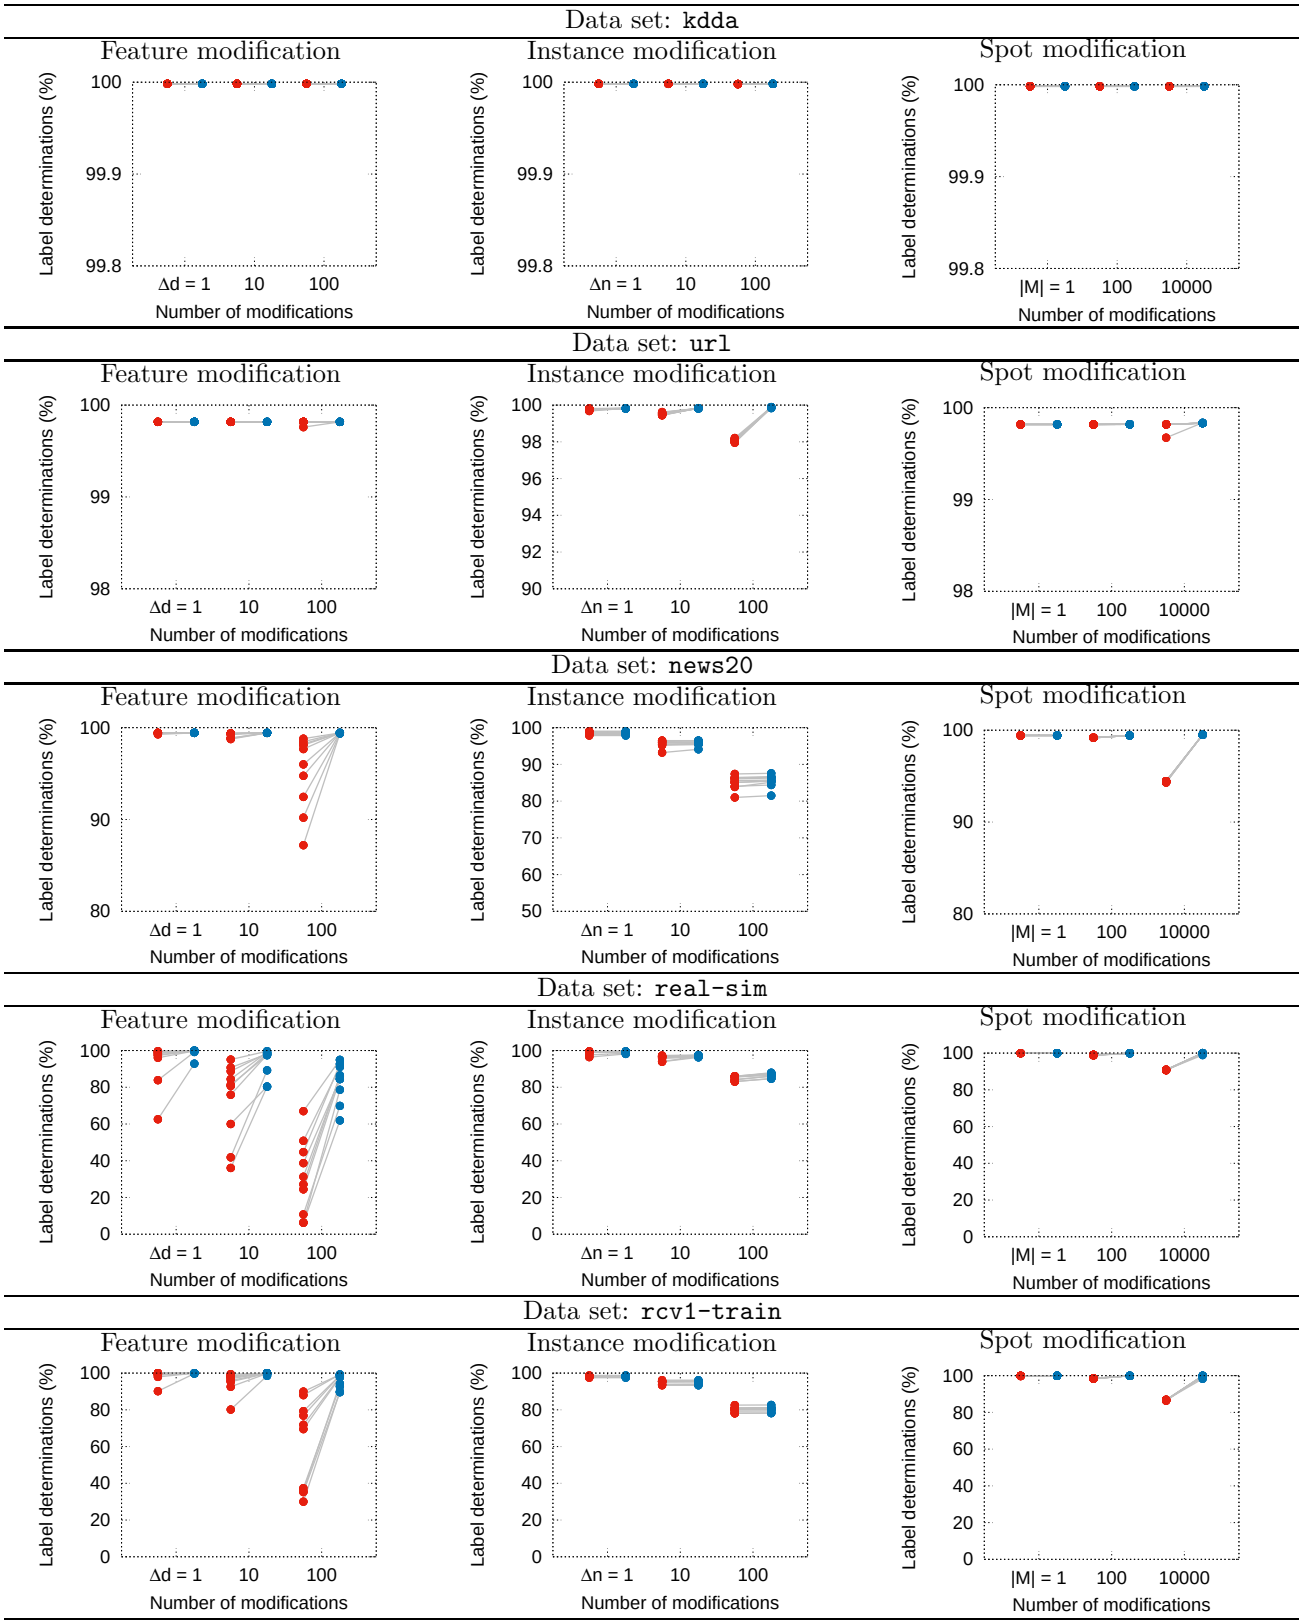

Figure 4: Label determination rate for $\lambda = 0.1$

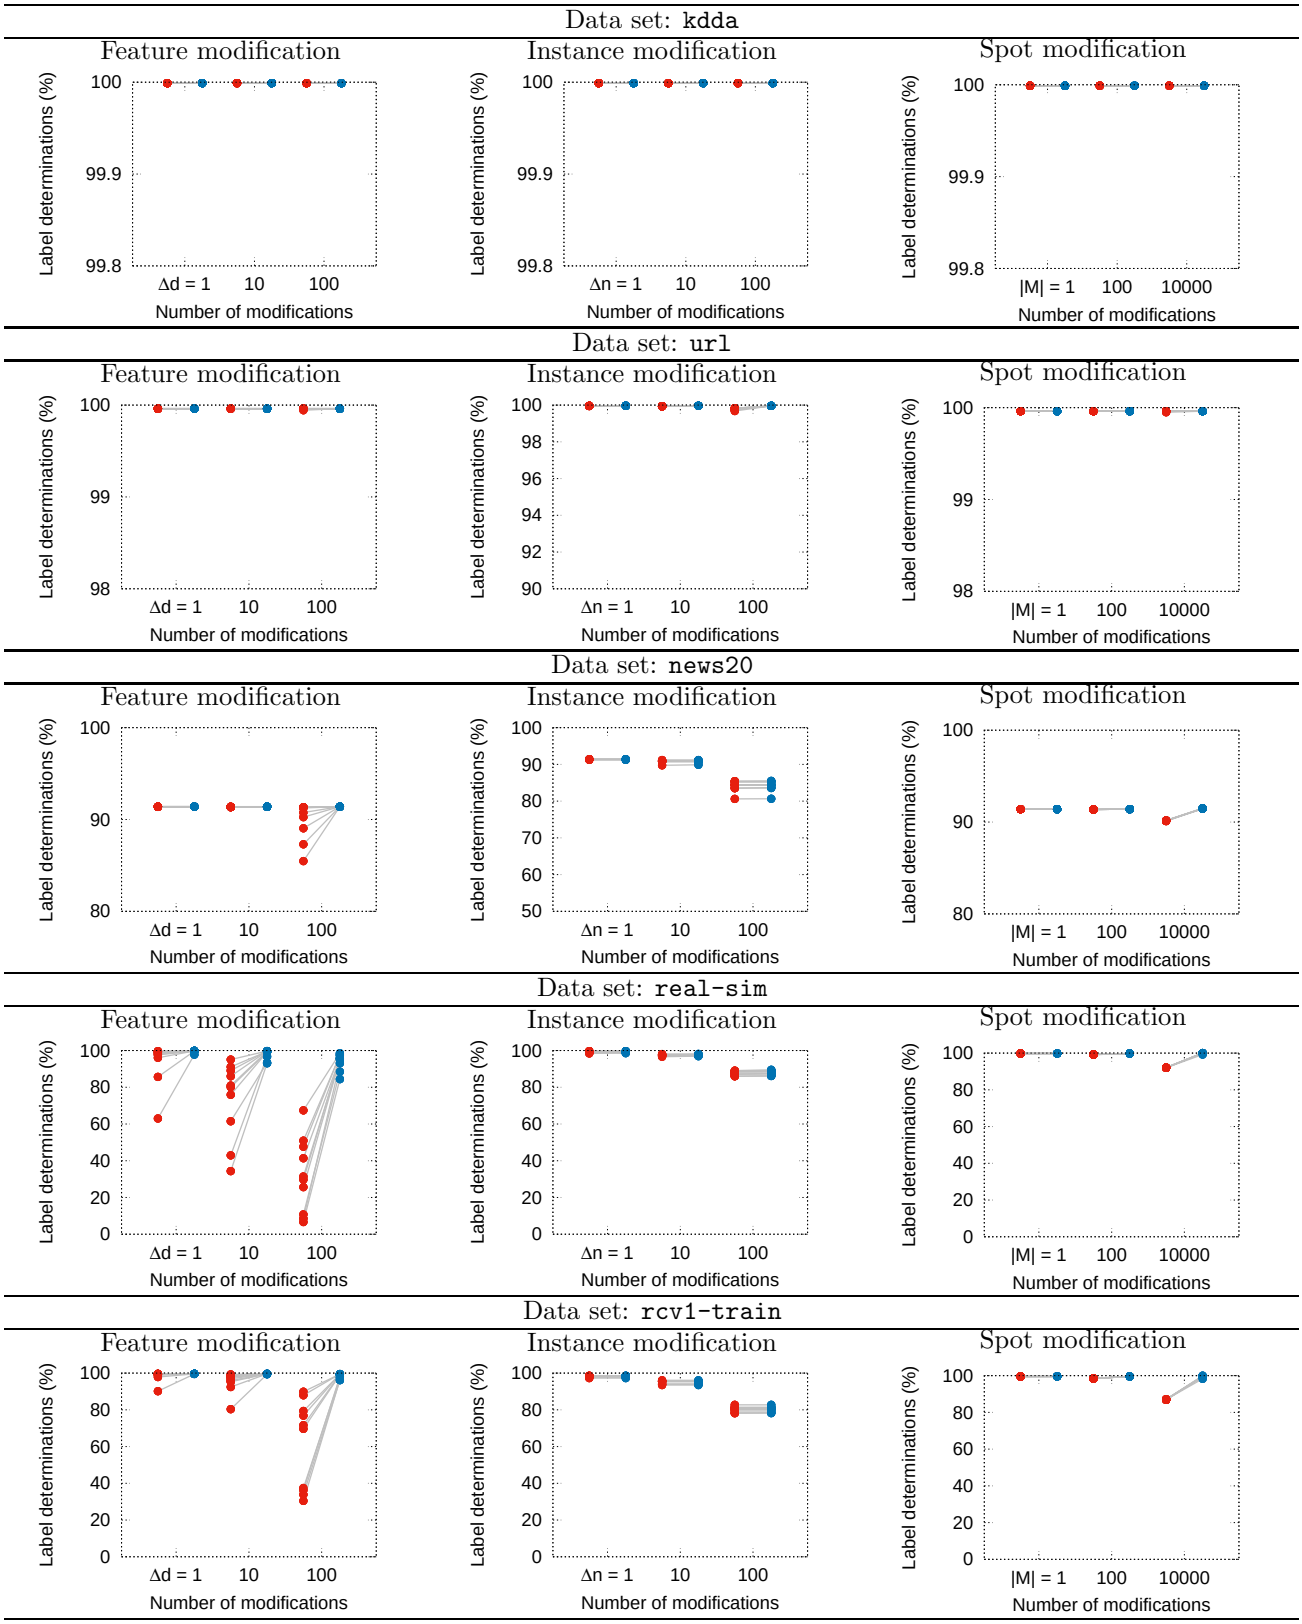

Figure 5: Label determination rate for $\lambda = 1.0$

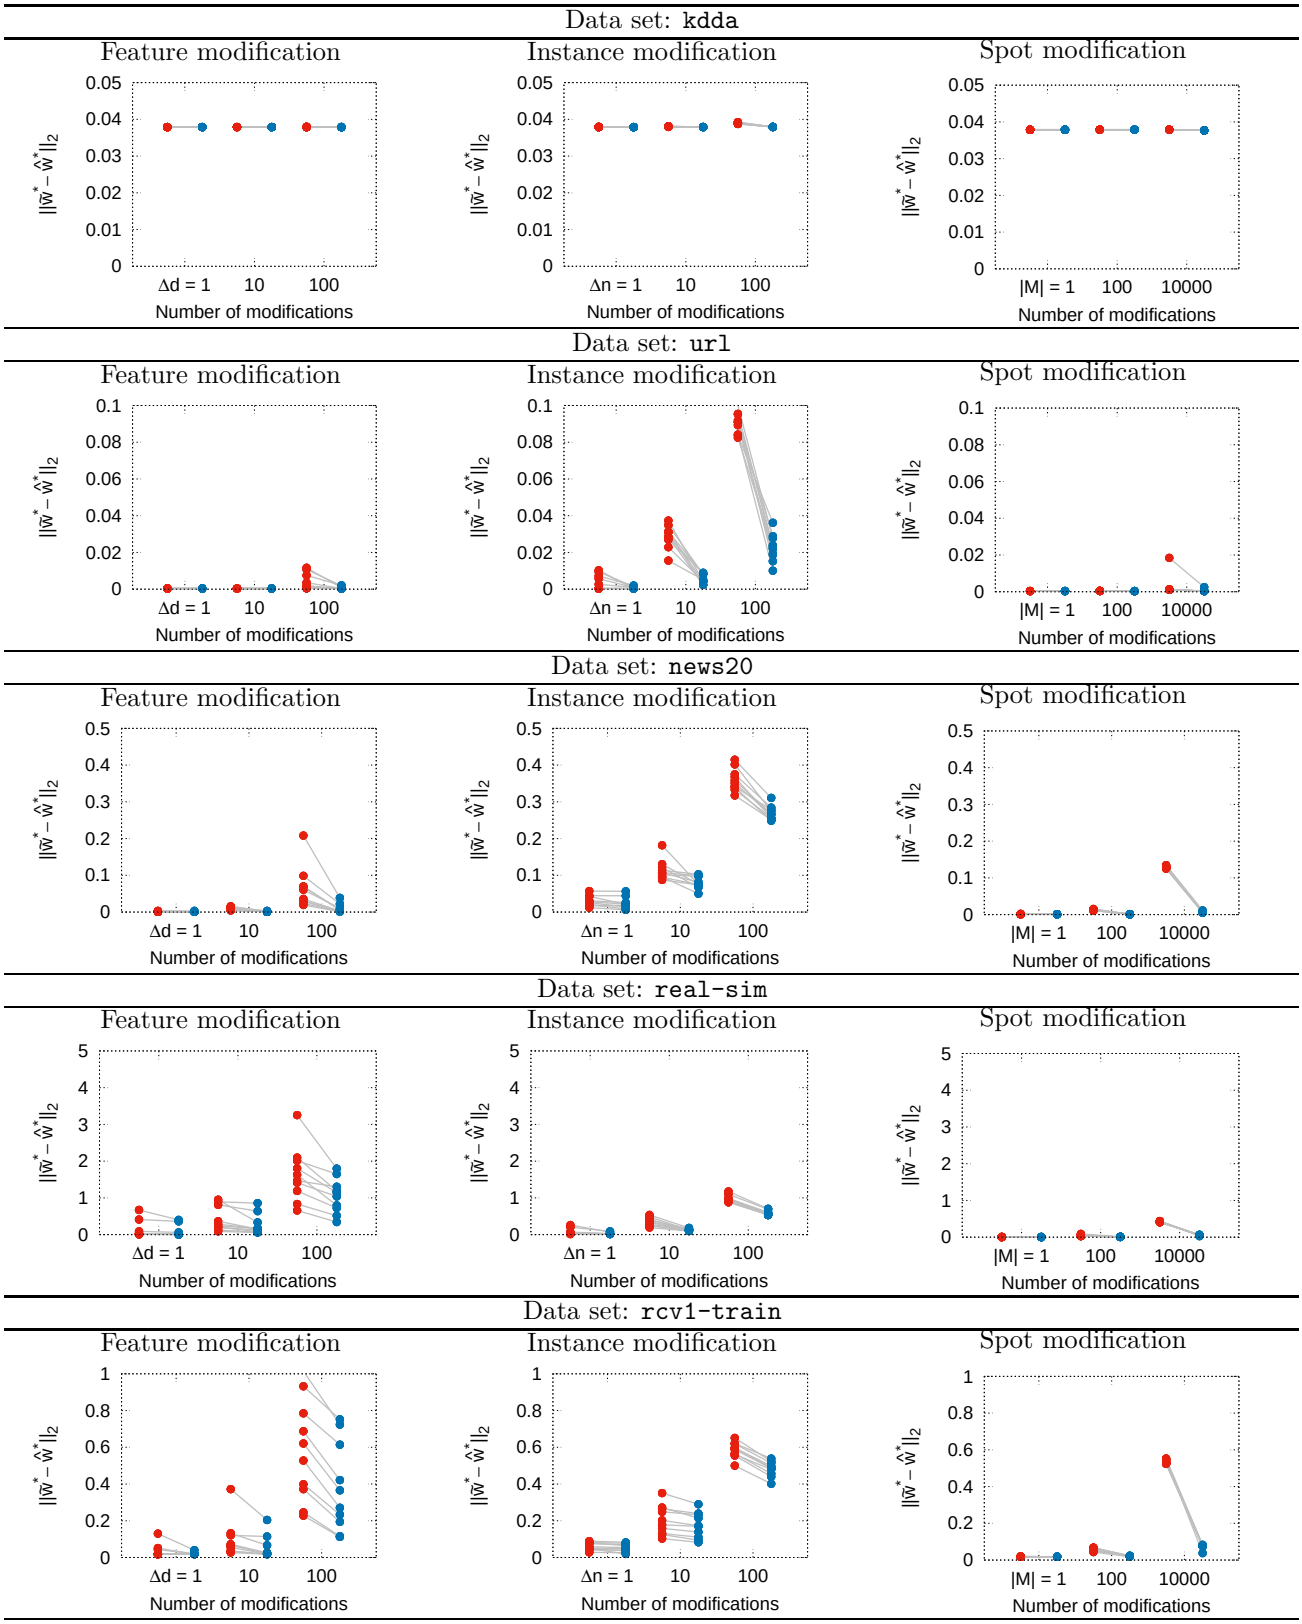

Figure 6: Upper bound of the parameter change $\|\tilde{\boldsymbol{w}}^* - \hat{\boldsymbol{w}}^*\|_2$ for $\lambda = 0.001$

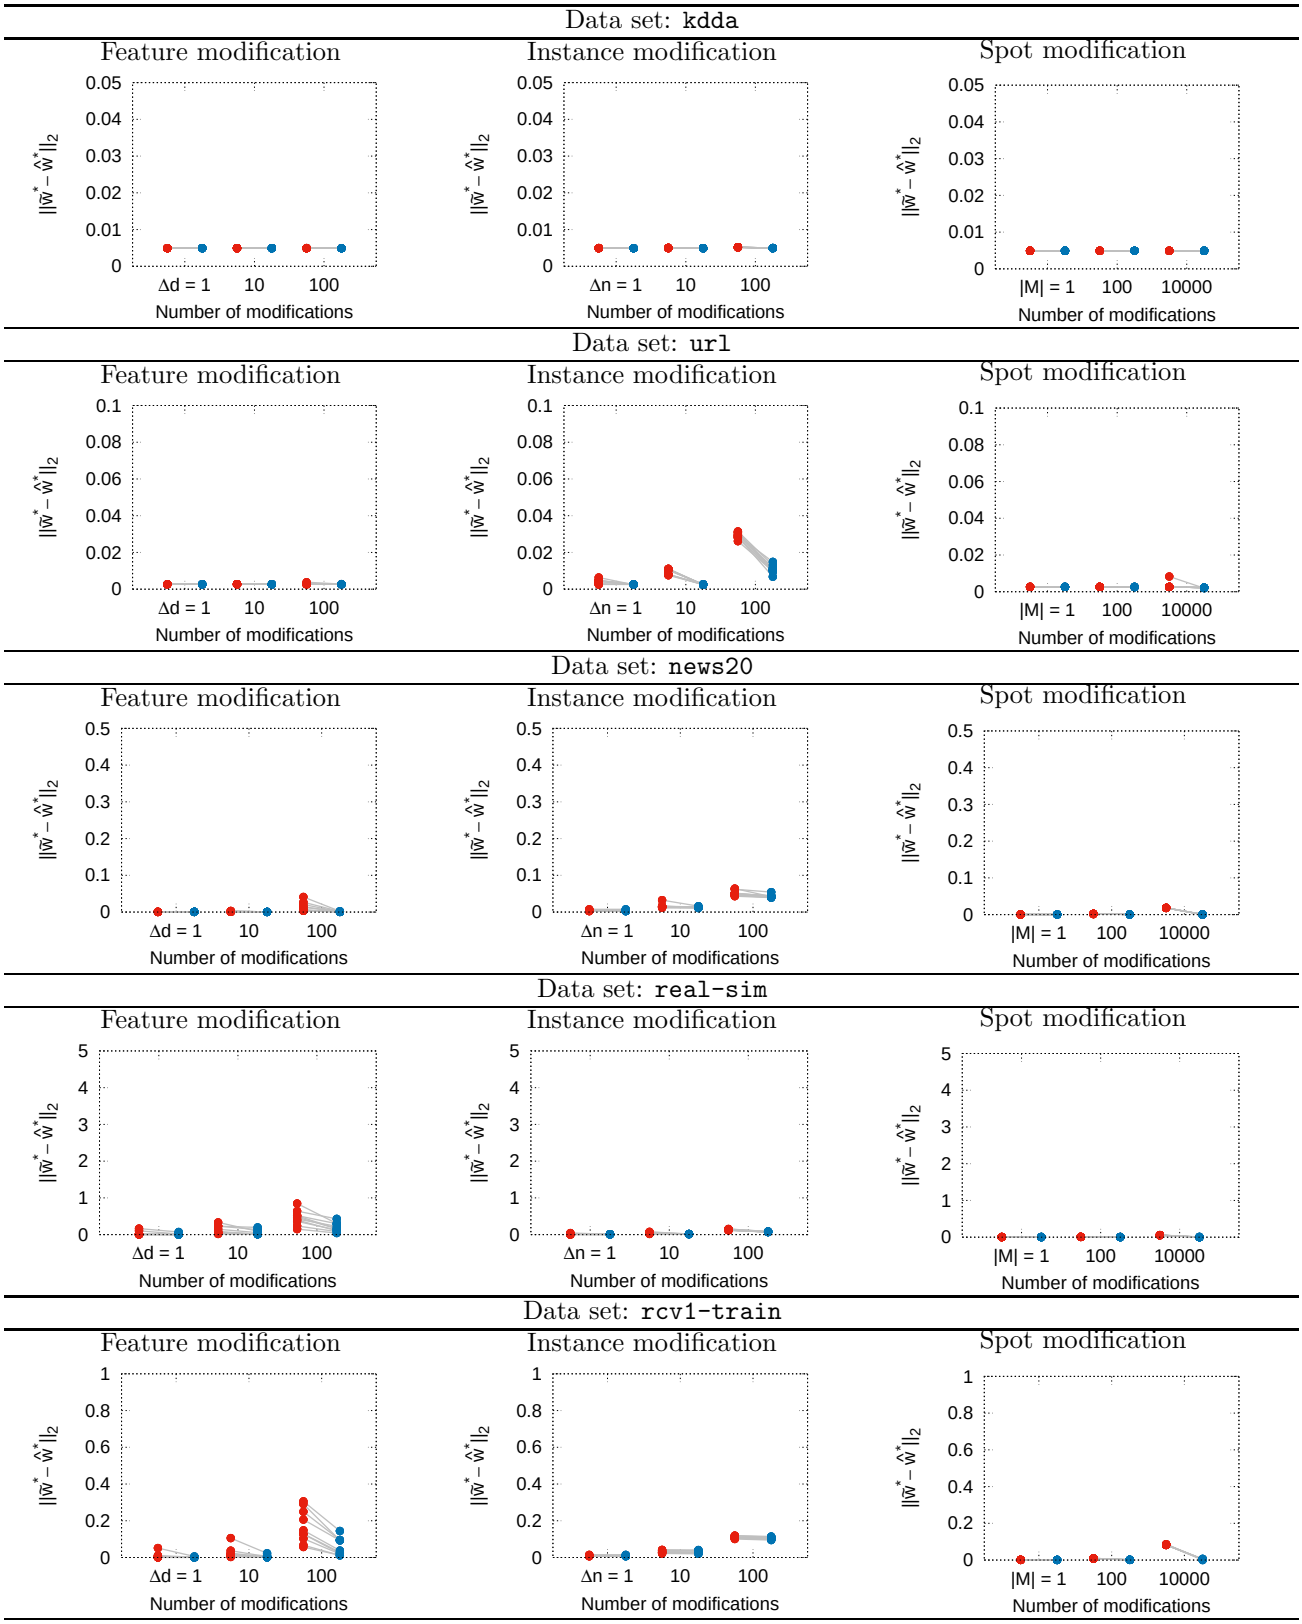

Figure 7: Upper bound of the parameter change $\|\tilde{\boldsymbol{w}}^* - \hat{\boldsymbol{w}}^*\|_2$ for $\lambda = 0.01$

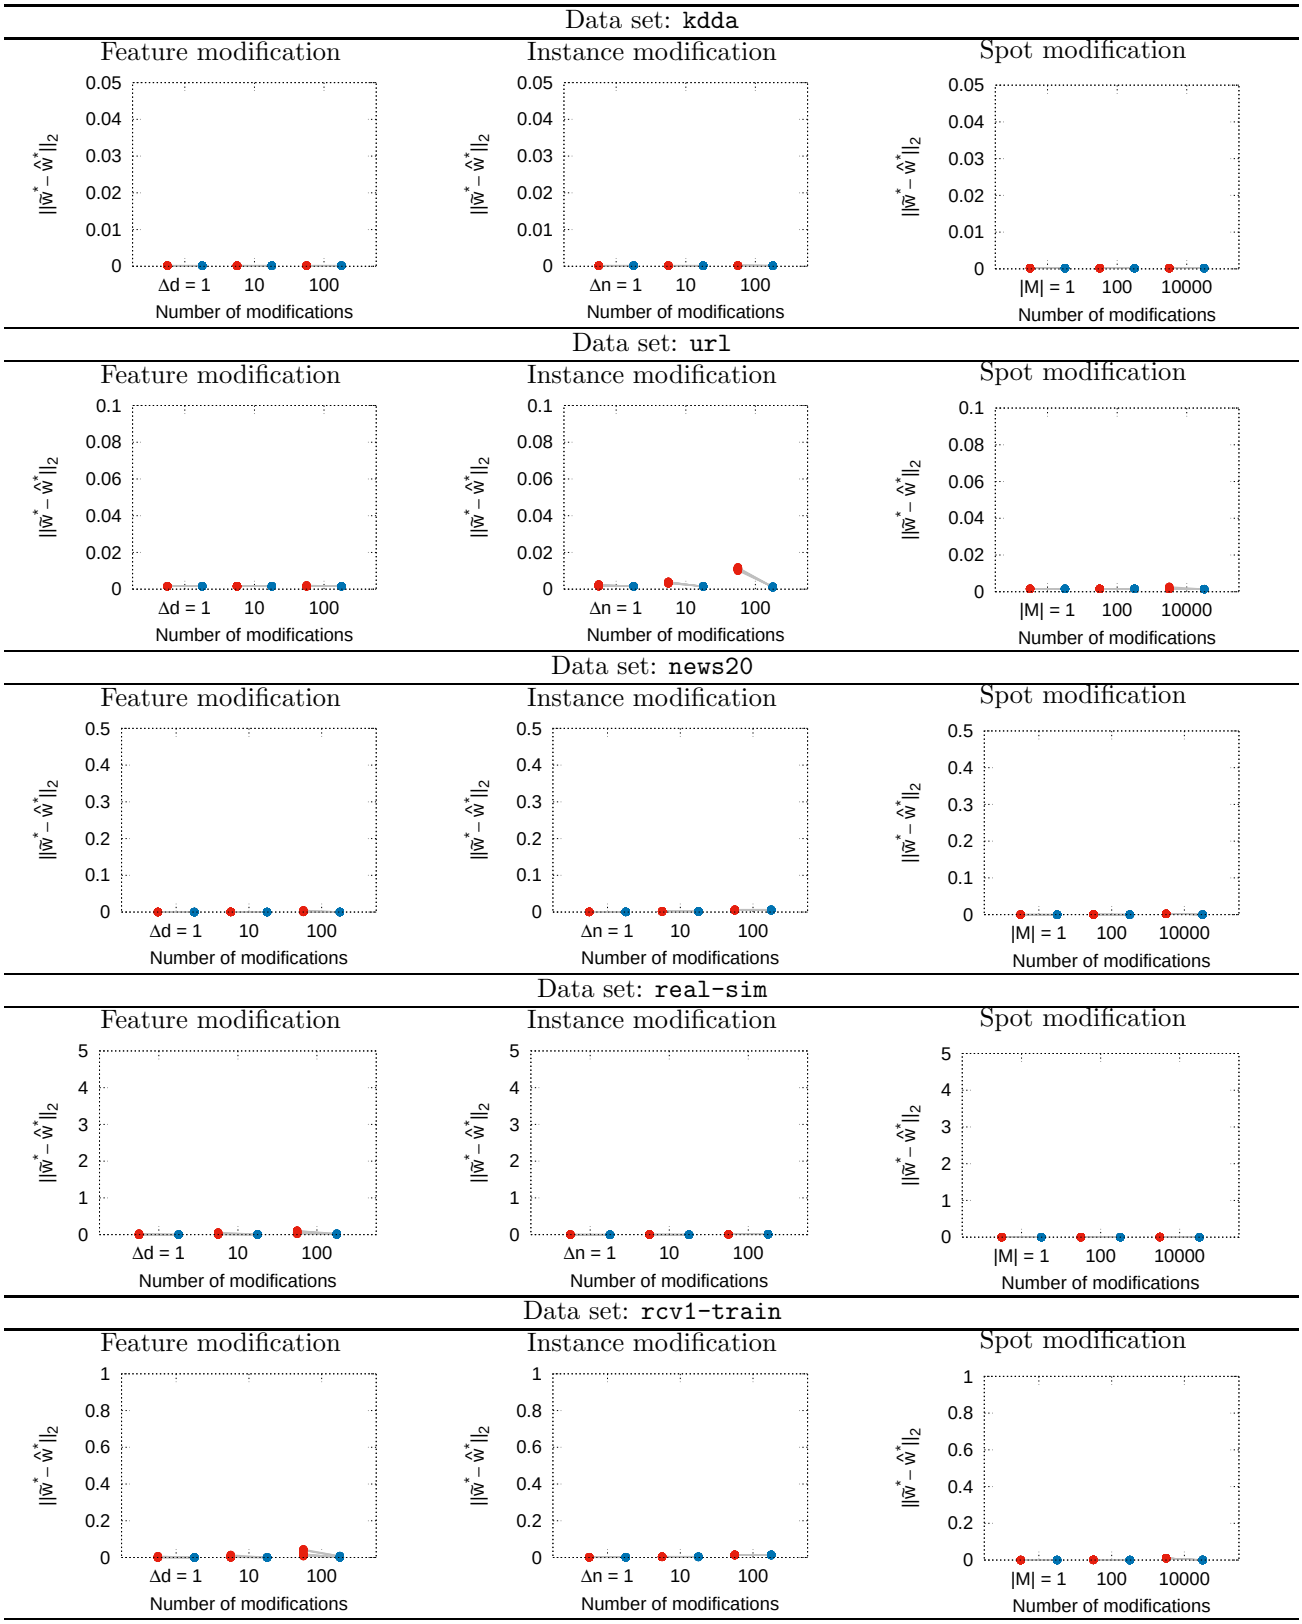

Figure 8: Upper bound of the parameter change $\|\tilde{\boldsymbol{w}}^* - \hat{\boldsymbol{w}}^*\|_2$ for $\lambda = 0.1$

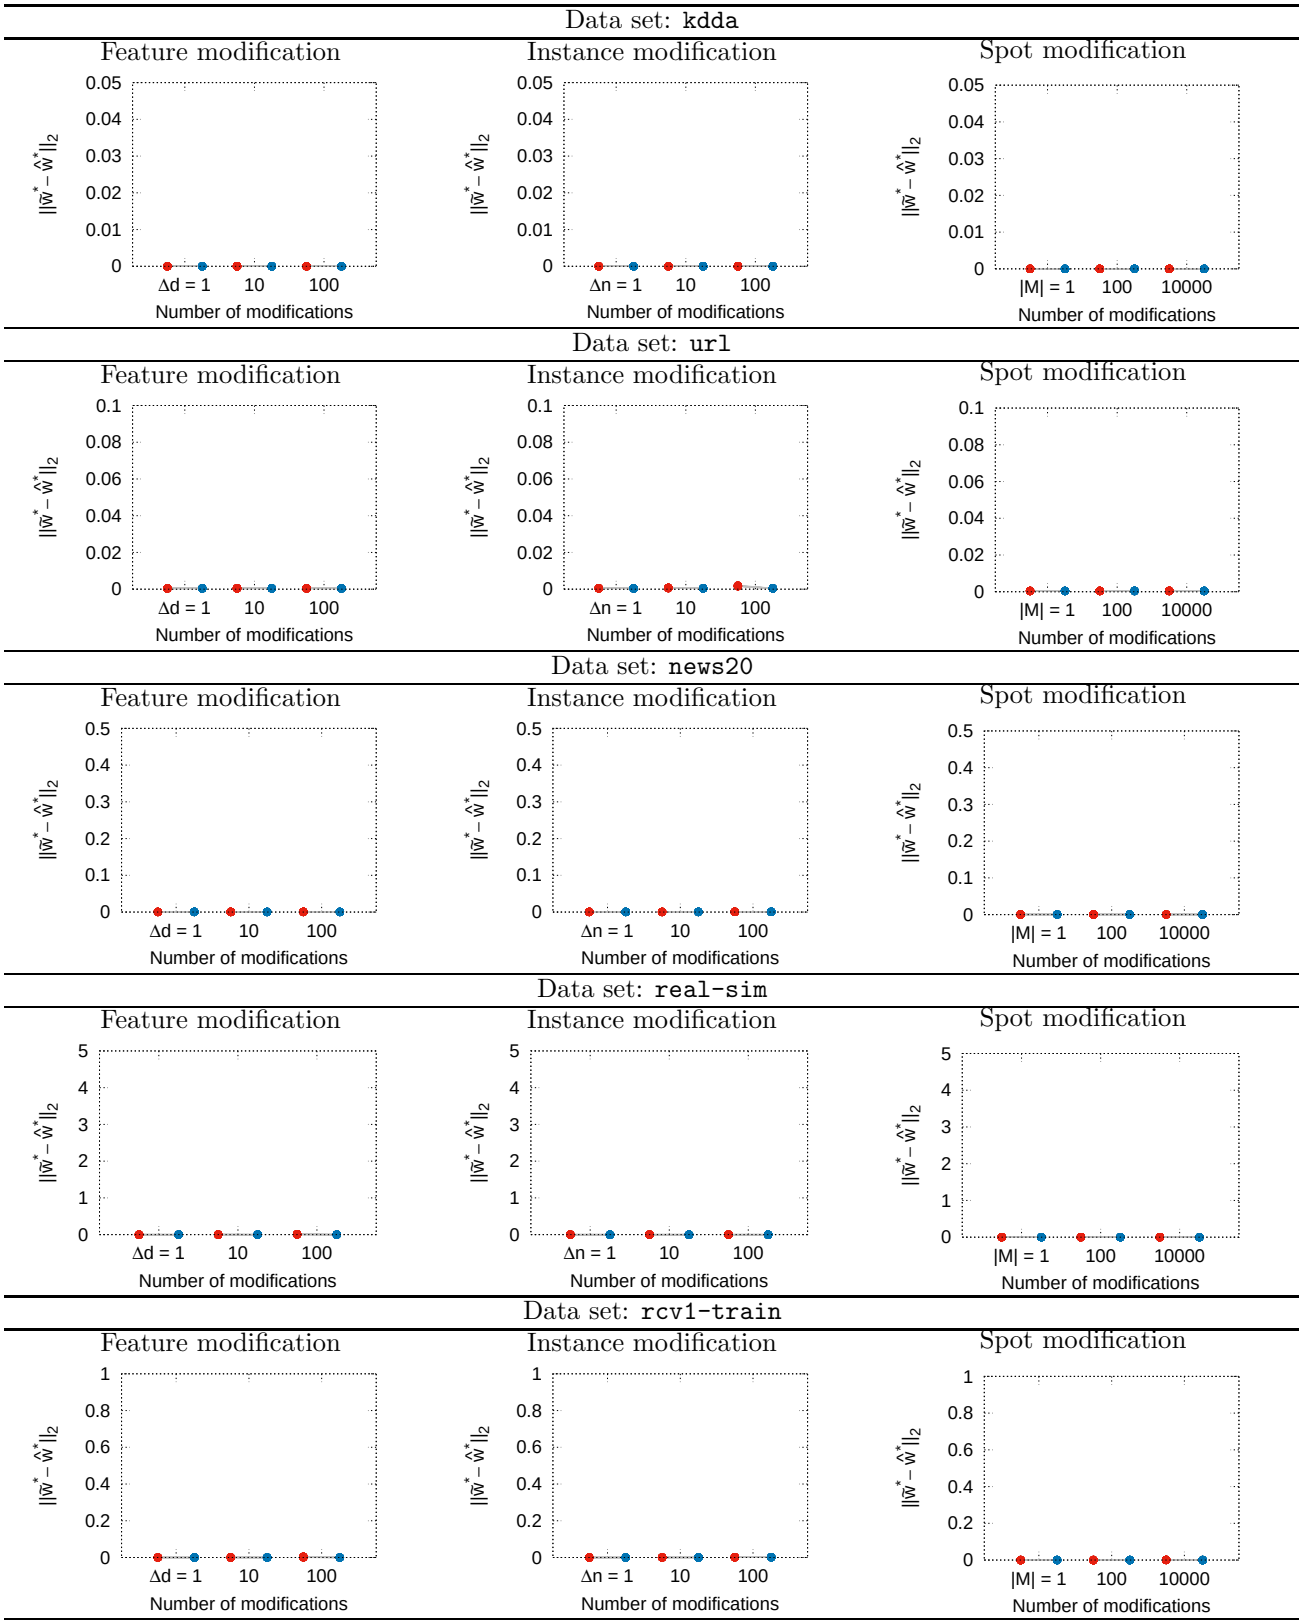

Figure 9: Upper bound of the parameter change $\|\tilde{\boldsymbol{w}}^* - \hat{\boldsymbol{w}}^*\|_2$ for $\lambda = 1.0$

Table 2: The ratio of the computation time of the bound computation to re-training (with warm-start). In all settings, the former cost is almost negligible.

	(a) spot modification			(b) instance modification			(c) feature modification		
	1	100	10000	1	10	100	1	10	100
				$\lambda = 01$					
kdd-a	3×10^{-6}	1×10^{-4}	1×10^{-4}	5×10^{-7}	2×10^{-4}	1×10^{-2}	2×10^{-7}	5×10^{-7}	5×10^{-5}
url	5×10^{-8}	2×10^{-5}	5×10^{-5}	3×10^{-7}	2×10^{-3}	3×10^{-2}	3×10^{-5}	3×10^{-6}	4×10^{-5}
news20	2×10^{-6}	3×10^{-5}	5×10^{-3}	7×10^{-5}	2×10^{-3}	6×10^{-2}	4×10^{-6}	1×10^{-5}	4×10^{-4}
real-sim	9×10^{-6}	1×10^{-4}	7×10^{-3}	5×10^{-5}	4×10^{-4}	2×10^{-3}	1×10^{-4}	6×10^{-4}	2×10^{-3}
rcv1-train	2×10^{-5}	2×10^{-4}	1×10^{-2}	2×10^{-4}	5×10^{-4}	6×10^{-3}	5×10^{-5}	2×10^{-4}	2×10^{-3}
				$\lambda = 0$					
kdd-a	6×10^{-6}	3×10^{-6}	2×10^{-4}	1×10^{-6}	1×10^{-4}	1×10^{-2}	5×10^{-7}	1×10^{-6}	2×10^{-4}
url	1×10^{-7}	3×10^{-6}	8×10^{-5}	6×10^{-6}	3×10^{-3}	5×10^{-2}	9×10^{-8}	1×10^{-6}	3×10^{-5}
news20	3×10^{-6}	6×10^{-5}	7×10^{-3}	8×10^{-4}	3×10^{-3}	5×10^{-2}	1×10^{-5}	4×10^{-5}	7×10^{-4}
real-sim	1×10^{-5}	2×10^{-4}	9×10^{-3}	7×10^{-5}	4×10^{-4}	3×10^{-3}	2×10^{-4}	1×10^{-3}	4×10^{-3}
rcv1-train	2×10^{-5}	3×10^{-4}	2×10^{-2}	2×10^{-4}	1×10^{-3}	9×10^{-3}	8×10^{-5}	4×10^{-4}	2×10^{-3}
				$\lambda = 0.1$					
kdd-a	2×10^{-7}	4×10^{-6}	1×10^{-4}	4×10^{-7}	5×10^{-6}	2×10^{-2}	2×10^{-7}	2×10^{-6}	1×10^{-5}
url	1×10^{-7}	2×10^{-6}	8×10^{-5}	4×10^{-7}	4×10^{-4}	5×10^{-2}	2×10^{-7}	1×10^{-6}	5×10^{-5}
news20	4×10^{-6}	4×10^{-5}	8×10^{-3}	6×10^{-5}	3×10^{-3}	1×10^{-2}	6×10^{-6}	3×10^{-5}	7×10^{-4}
real-sim	1×10^{-5}	2×10^{-4}	2×10^{-2}	8×10^{-5}	5×10^{-4}	5×10^{-3}	2×10^{-4}	1×10^{-3}	6×10^{-3}
rcv1-train	3×10^{-5}	4×10^{-4}	2×10^{-2}	3×10^{-4}	2×10^{-3}	9×10^{-3}	9×10^{-5}	5×10^{-4}	1×10^{-3}
				$\lambda = 1$					
kdd-a	3×10^{-7}	3×10^{-6}	2×10^{-4}	5×10^{-7}	6×10^{-6}	2×10^{-2}	2×10^{-7}	2×10^{-6}	2×10^{-5}
url	1×10^{-7}	2×10^{-6}	3×10^{-4}	5×10^{-7}	8×10^{-3}	5×10^{-2}	2×10^{-7}	2×10^{-6}	7×10^{-5}
news20	4×10^{-6}	1×10^{-4}	6×10^{-3}	2×10^{-4}	2×10^{-3}	9×10^{-3}	6×10^{-6}	2×10^{-5}	1×10^{-3}
real-sim	3×10^{-5}	2×10^{-4}	1×10^{-2}	1×10^{-4}	5×10^{-4}	5×10^{-3}	2×10^{-4}	2×10^{-3}	8×10^{-3}
rcv1-train	5×10^{-5}	4×10^{-4}	2×10^{-2}	3×10^{-4}	2×10^{-3}	2×10^{-2}	2×10^{-4}	4×10^{-4}	4×10^{-3}

6 Conclusions

In this paper, we present a method for quickly incorporating the data modification effect into the classifier. The proposed method provides bounds on the optimal solution with the cost proportional to the size of the data modification. The experimental results indicate that our bound computation method in §3 is highly effective when the number of modified elements is much smaller than the entire dataset size. In addition, partial optimization approach is also effective especially when the bounds in Theorem 1 is not good enough.

References

[1] D. DeCoste and K. Wagstaff. Alpha seeding for support vector machines. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2000.

[2] G. Cauwenberghs and T. Poggio. Incremental and decremental support vector machine learning. In *Advances in Neural Information Processing Systems*, 2001.

[3] Laurent El Ghaoui, Vivian Viallon, and Tarek Rabbani. Safe feature elimination for the lasso and sparse supervised learning problems. *Pacific Journal of Optimization*, 8(4):667–698, 2012.

[4] Zhen J Xiang, Hao Xu, and Peter J Ramadge. Learning sparse representations of high dimensional data on large scale dictionaries. In *Advances in Neural Information Processing Systems*, pages 900–908, 2011.

[5] Jie Wang, Jiayu Zhou, Peter Wonka, and Jieping Ye. Lasso screening rules via dual polytope projection. In *Advances in Neural Information Processing Systems*, pages 1070–1078, 2013.

[6] Antoine Bonnefoy, Valentin Emiya, Liva Ralaivola, and Rémi Gribonval. A dynamic screening principle for the lasso. In *Signal Processing Conference (EUSIPCO), 2014 Proceedings of the 22nd European*, pages 6–10. IEEE, 2014.

[7] Jun Liu, Zheng Zhao, Jie Wang, and Jieping Ye. Safe Screening with Variational Inequalities and Its Application to Lasso. In *Proceedings of the 31st International Conference on Machine Learning*, 2014.

[8] Jie Wang, Jiayu Zhou, Jun Liu, Peter Wonka, and Jieping Ye. A safe screening rule for sparse logistic regression. In *Advances in Neural Information Processing Systems*, pages 1053–1061, 2014.

[9] Zhen James Xiang, Yun Wang, and Peter J Ramadge. Screening tests for lasso problems. *arXiv preprint arXiv:1405.4897*, 2014.

[10] Olivier Fercoq, Alexandre Gramfort, and Joseph Salmon. Mind the duality gap: safer rules for the lasso. In *Proceedings of the 32nd International Conference on Machine Learning*, pages 333–342, 2015.

[11] Eugene Ndiaye, Olivier Fercoq, Alexandre Gramfort, and Joseph Salmon. Gap safe screening rules for sparse multi-task and multi-class models. In *Advances in Neural Information Processing Systems*, pages 811–819, 2015.

[12] Shota Okumura, Yoshiki Suzuki, and Ichiro Takeuchi. Quick sensitivity analysis for incremental data modification and its application to leave-one-out cv in linear classification problems. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 885–894. ACM, 2015.

[13] Moshe Gabel, Daniel Keren, and Assaf Schuster. Monitoring least squares models of distributed streams. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 319–328. ACM, 2015.

[14] Ralph Tyrell Rockafellar. *Convex analysis*. Princeton university press, 1970.

[15] Atsushi Shibagaki, Yoshiki Suzuki, Masayuki Karasuyama, and Ichiro Takeuchi. Regularization path of cross-validation error lower bounds. In *Advances in Neural Information Processing Systems*, pages 1666–1674, 2015.

[16] Kohei Ogawa, Yoshiki Suzuki, and Ichiro Takeuchi. Safe screening of non-support vectors in pathwise svm computation. In *Proceedings of the 30th International Conference on Machine Learning*, pages 1382–1390, 2013.

[17] Jie Wang, Peter Wonka, and Jieping Ye. Scaling svm and least absolute deviations via exact data reduction. *Proceedings of The 31st International Conference on Machine Learning*, 2014.

[18] Julian Zimmert, Christian Schroeder de Witt, Giancarlo Kerg, and Marius Kloft. Safe screening for support vector machines. *NIPS 2015 Workshop on Optimization in Machine Learning (OPT)*, 2015.

[19] Chih-Chung Chang and Chih-Jen Lin. Libsvm: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2(3):27, 2011.

Appendix

A Proofs

Before showing the proofs, we state some lemmas used in the proofs.

Lemma 1. *(Theorem 2 in [11]) Let $\hat{\boldsymbol{w}}$ be an any solution of the primal problem P, and $\hat{\boldsymbol{\alpha}}$ be an any feasible solution of the dual problem D. If D is γ-strongly concave then the optimal solution of dual problem $\boldsymbol{\alpha}^*$ is within in the sphere $\Theta_{\boldsymbol{\alpha}^*} := \{\boldsymbol{\alpha} \mid \|\boldsymbol{\alpha} - \hat{\boldsymbol{\alpha}}\|_2 \leq \sqrt{2\gamma^{-1}(P(\hat{\boldsymbol{w}}) - D(\hat{\boldsymbol{\alpha}}))}\}$.*

Lemma 2. *(Corollary 4.2 in [18]) Let $\hat{\boldsymbol{w}}$ be an any solution of the primal problem P, and $\hat{\boldsymbol{\alpha}}$ be an any feasible solution of the dual problem D. If P is λ-strongly convex then the optimal solution of primal problem $\boldsymbol{w}^*$ is within in the sphere $\Theta_{\boldsymbol{w}^*} := \{\boldsymbol{w} \mid \|\boldsymbol{w} - \hat{\boldsymbol{w}}\|_2 \leq \sqrt{2\lambda^{-1}(P(\hat{\boldsymbol{w}}) - D(\hat{\boldsymbol{\alpha}}))}\}$.*

Lemma 3. *Let $\boldsymbol{\eta}, \boldsymbol{c} \in \mathbb{R}^d$ be arbitrary vectors and a scalar $p > 0$. Then, the following optimization problems have closed form solutions as follows:*

$$\boldsymbol{\eta}^\top \boldsymbol{c} - \sqrt{p}\|\boldsymbol{\eta}\|_2 = \min_{\boldsymbol{q} \in \mathbb{R}^d} \boldsymbol{\eta}^\top \boldsymbol{q} \quad \text{s.t.} \quad \|\boldsymbol{q} - \boldsymbol{c}\|_2 \leq \sqrt{p},$$
$$\boldsymbol{\eta}^\top \boldsymbol{c} + \sqrt{p}\|\boldsymbol{\eta}\|_2 = \max_{\boldsymbol{q} \in \mathbb{R}^d} \boldsymbol{\eta}^\top \boldsymbol{q} \quad \text{s.t.} \quad \|\boldsymbol{q} - \boldsymbol{c}\|_2 \leq \sqrt{p}$$

Proof. Using a Lagrange multiplier method, we can easily prove this Lemma. □

A.1 Proof of Theorem 1

Proof. From the assumptions with regard to the penalty function ψ,

$$\begin{aligned}
\psi^*(\boldsymbol{w}) &= \sup_{\boldsymbol{v} \in \mathbb{R}^d} \left\{ \boldsymbol{w}^\top \boldsymbol{v} - \psi(\boldsymbol{v}) \right\} \\
&= \sup_{\boldsymbol{v} \in \mathbb{R}^d} \left\{ \sum_{j \in [d]} w_j v_j - \sum_{j \in [d]} \psi_j(v_j) \right\} \\
&= \sum_{j \in [d]} \left\{ \sup_{v_j \in \mathbb{R}} \{ w_j v_j - \psi_j(v_j) \} \right\}.
\end{aligned}$$

Let $\psi_j^*(w_j)$ be a function such that $\sup_{v_j \in \mathbb{R}} \{w_j v_j - \psi_j(v_j)\}$, $\psi^*(\boldsymbol{w}) = \sum_{j \in [d]} \psi_j^*(w_j)$. Thus, from KKT condition,

$$\tilde{w}_j^* \in \partial \psi_j^*(n^{-1} \tilde{\boldsymbol{z}}_{\cdot j}^\top \tilde{\boldsymbol{\alpha}}^*). \tag{9}$$

The old and the new primal and dual problems are respectively denoted as $\hat{P}, \tilde{P}, \hat{D}, \tilde{D}$. Since $\hat{\boldsymbol{w}}^*$ and $\hat{\boldsymbol{\alpha}}^*$ is optimal solution of $\hat{P}$ and $\hat{D}$, respectively,

$$\begin{aligned}\tilde{P}(\hat{\boldsymbol{w}}^*) - \tilde{D}(\hat{\boldsymbol{\alpha}}^*) &= n^{-1}\sum_{i\in[n]} \phi(\hat{\boldsymbol{z}}_{i\cdot}^\top \hat{\boldsymbol{w}}^*) + n^{-1}\sum_{i\in\mathcal{M}_i} (\phi(\tilde{\boldsymbol{z}}_{i\cdot}^\top \hat{\boldsymbol{w}}^*) - \phi(\hat{\boldsymbol{z}}_{i\cdot}^\top \hat{\boldsymbol{w}}^*)) + \psi(\hat{\boldsymbol{w}}^*) \\ &\quad + n^{-1}\sum_{i\in[n]} \phi^*(-\alpha_i^*) + \sum_{j\in[d]} (\psi_j^*(n^{-1}\hat{\boldsymbol{z}}_{\cdot j}^\top \hat{\boldsymbol{\alpha}}^*)) + \sum_{j\in\mathcal{M}_j} (\psi_j^*(n^{-1}\tilde{\boldsymbol{z}}_{\cdot j}^\top \hat{\boldsymbol{\alpha}}^*) - \psi_j^*(n^{-1}\hat{\boldsymbol{z}}_{\cdot j}^\top \hat{\boldsymbol{\alpha}}^*)) \\ &= n^{-1}\sum_{i\in\mathcal{M}_i} (\phi(\tilde{\boldsymbol{z}}_{i\cdot}^\top \hat{\boldsymbol{w}}^*) - \phi(\hat{\boldsymbol{z}}_{i\cdot}^\top \hat{\boldsymbol{w}}^*)) + \sum_{j\in\mathcal{M}_j} (\psi_j^*(n^{-1}\tilde{\boldsymbol{z}}_{\cdot j}^\top \hat{\boldsymbol{\alpha}}^*) - \psi_j^*(n^{-1}\hat{\boldsymbol{z}}_{\cdot j}^\top \hat{\boldsymbol{\alpha}}^*)). \end{aligned} \tag{10}$$

For a subset $\mathcal{J} \subseteq [d]$, we define $\mathcal{M}(\cdot, \mathcal{J}) := \{i \in [n] \mid \exists j \in \mathcal{J} \text{ s.t. } (i,j) \in \mathcal{M}\}$. Similarly, for a subset $\mathcal{I} \subseteq [n]$, we define $\mathcal{M}(\mathcal{I}, \cdot) := \{j \in [n] \mid \exists i \in \mathcal{I} \text{ s.t. } (i,j) \in \mathcal{M}\}$. Since $\{\hat{\boldsymbol{z}}_{i\cdot}^\top \hat{\boldsymbol{w}}^*\}_{i\in[n]}$ are stored in the memory and $\tilde{\boldsymbol{z}}_{i\cdot}^\top \hat{\boldsymbol{w}}^* = \hat{\boldsymbol{z}}_{i\cdot}^\top \hat{\boldsymbol{w}}^* + \sum_{j\in\mathcal{M}(i,\cdot)} \hat{w}_j^*(\tilde{z}_{ij} - \hat{z}_{ij})$, $\{\tilde{\boldsymbol{z}}_{i\cdot}^\top \hat{\boldsymbol{w}}^*\}_{i\in\mathcal{M}_i}$ is computed with $\mathcal{O}(\sum_{i\in\mathcal{M}_i} |\mathcal{M}(i,\cdot)|) = \mathcal{O}(|\mathcal{M}|)$ time complexity. Similarly, $\{\tilde{\boldsymbol{z}}_{\cdot j}^\top \hat{\boldsymbol{w}}^*\}_{j\in\mathcal{M}_j}$ is computed with $\mathcal{O}(\sum_{j\in\mathcal{M}_j} |\mathcal{M}(\cdot,j)|) = \mathcal{O}(|\mathcal{M}|)$ time complexity because $\tilde{\boldsymbol{z}}_{\cdot j}^\top \hat{\boldsymbol{w}}^* = \hat{\boldsymbol{z}}_{\cdot j}^\top \hat{\boldsymbol{\alpha}}^* + \sum_{i\in\mathcal{M}(\cdot,j)} \hat{\alpha}_i^*(\tilde{z}_{ij} - \hat{z}_{ij})$. Also, the computational cost of $\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)$ is $\mathcal{O}(|\mathcal{M}|)$.

(i) The case where the loss function has **Property A**:

From **Property A**, ϕ^* is γ-strongly convex and so $\tilde{D}$ is γ/n-strongly concave and therefore, $\tilde{\boldsymbol{\alpha}}^* \in \Theta_{\tilde{\boldsymbol{\alpha}}^*} := \{\boldsymbol{\alpha} \mid \|\boldsymbol{\alpha} - \hat{\boldsymbol{\alpha}}^*\|_2 \leq \sqrt{2\gamma^{-1}(\tilde{P}(\hat{\boldsymbol{w}}^*) - \tilde{D}(\hat{\boldsymbol{\alpha}}^*))}\}$ by using Lemma 1. Since $\Theta_{\tilde{\boldsymbol{\alpha}}^*}$ is a sphere, the lower and upper bounds of the inner product of arbitrary vector $\boldsymbol{\eta} \in \mathbb{R}^n$ and $\tilde{\boldsymbol{\alpha}}^*$ are given in closed form as follows by using Lemma 3,

$$\boldsymbol{\eta}^\top \tilde{\boldsymbol{\alpha}}^* \geq \boldsymbol{\eta}^\top \hat{\boldsymbol{\alpha}}^* - \|\boldsymbol{\eta}\|_2 \sqrt{2n\gamma^{-1}(\tilde{P}(\hat{\boldsymbol{w}}^*) - \tilde{D}(\hat{\boldsymbol{\alpha}}^*))}, \tag{11}$$

$$\boldsymbol{\eta}^\top \tilde{\boldsymbol{\alpha}}^* \leq \boldsymbol{\eta}^\top \hat{\boldsymbol{\alpha}}^* + \|\boldsymbol{\eta}\|_2 \sqrt{2n\gamma^{-1}(\tilde{P}(\hat{\boldsymbol{w}}^*) - \tilde{D}(\hat{\boldsymbol{\alpha}}^*))}. \tag{12}$$

From (9), (10), (11), the assumptions and since $\partial\psi_j^*$ is a monotonically increasing function because ψ_j^* is a convex function, $\tilde{w}_j^*$ is bounded as by using Lemma 3

$$\begin{aligned}\tilde{w}_j^* &\geq \inf \partial\psi_j^*(n^{-1}\tilde{\boldsymbol{z}}_{\cdot j}^\top \tilde{\boldsymbol{\alpha}}^*) \\ &\geq \inf \partial\psi_j^* \left(\max \left\{ \inf_{\boldsymbol{\alpha}\in \mathrm{dom}\tilde{D}(\boldsymbol{\alpha})} n^{-1}\tilde{\boldsymbol{z}}_{\cdot j}^\top \boldsymbol{\alpha}, \quad \min_{\boldsymbol{\alpha}\in\Theta_{\tilde{\boldsymbol{\alpha}}^*}} n^{-1}\tilde{\boldsymbol{z}}_{\cdot j}^\top \boldsymbol{\alpha} \right\} \right) \\ &\geq \inf \partial\psi_j^* \left(\max \left\{ \inf_{\boldsymbol{\alpha}\in \mathrm{dom}\tilde{D}(\boldsymbol{\alpha})} n^{-1}\tilde{\boldsymbol{z}}_{\cdot j}^\top \boldsymbol{\alpha}, \quad n^{-1}\tilde{\boldsymbol{z}}_{\cdot j}^\top \hat{\boldsymbol{\alpha}}^* - \|\tilde{\boldsymbol{x}}_{\cdot j}\|_2 \sqrt{2\gamma^{-1}\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)} \right\} \right). \end{aligned}$$

Thus, $\tilde{w}_j^* \geq L_D(w_j^*)$. Similarly, from (9), (10), (12), the assumptions and since $\partial\psi_j^*$ is a monotonically

increasing function, $\tilde{w}_j^*$ is bounded as:

$$\begin{aligned}
\tilde{w}_j^* &\leq \sup \partial \psi_j^*(n^{-1}\tilde{\boldsymbol{z}}_{\cdot j}^\top \tilde{\boldsymbol{\alpha}}^*) \\
&\leq \sup \partial \psi_j^* \left(\min \left\{ \sup_{\boldsymbol{\alpha} \in \mathrm{dom}\tilde{D}(\boldsymbol{\alpha})} n^{-1}\tilde{\boldsymbol{z}}_{\cdot j}^\top \boldsymbol{\alpha}, \quad \max_{\boldsymbol{\alpha} \in \Theta_{\tilde{\boldsymbol{\alpha}}^*}} n^{-1}\tilde{\boldsymbol{z}}_{\cdot j}^\top \boldsymbol{\alpha} \right\} \right) \\
&\leq \sup \partial \psi_j^* \left(\min \left\{ \sup_{\boldsymbol{\alpha} \in \mathrm{dom}\tilde{D}(\boldsymbol{\alpha})} n^{-1}\tilde{\boldsymbol{z}}_{\cdot j}^\top \boldsymbol{\alpha}, \quad n^{-1}\tilde{\boldsymbol{z}}_{\cdot j}^\top \hat{\boldsymbol{\alpha}}^* + \|\tilde{\boldsymbol{x}}_{\cdot j}\|_2 \sqrt{2\gamma^{-1}\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)} \right\} \right).
\end{aligned}$$

Thus, $\tilde{w}_j^* \leq U_D(w_j^*)$. Then we bound $\tilde{\alpha}_i^*$. From Lemma 1 we have

$$\hat{\alpha}_i^* - \sqrt{2n\gamma^{-1}\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)} \leq \tilde{\alpha}_i^* \leq \hat{\alpha}_i^* + \sqrt{2n\gamma^{-1}\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)},$$

and thus

$$\begin{aligned}
\tilde{\alpha}_i^* &\geq \max \left\{ \inf_{\alpha \in \mathrm{dom}\tilde{D}(\alpha)} \alpha_i, \quad \hat{\alpha}_i^* - \sqrt{2n\gamma^{-1}\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)} \right\}, \\
\tilde{\alpha}_i^* &\leq \min \left\{ \sup_{\alpha \in \mathrm{dom}\tilde{D}(\alpha)} \alpha_i, \quad \hat{\alpha}_i^* + \sqrt{2n\gamma^{-1}\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)} \right\}.
\end{aligned}$$

The all quantities required for evaluating $L_D(\tilde{w}_j^*), U_D(\tilde{w}_j^*), L_D(\tilde{\alpha}_i^*)$ and $U_D(\tilde{\alpha}_i^*)$ are stored in the memory without $\|\tilde{\boldsymbol{x}}_{\cdot j}\|_2$ and $\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)$. Since $\|\tilde{\boldsymbol{x}}_{\cdot j}\|_2$ is stored in the memory and $\|\tilde{\boldsymbol{x}}_{\cdot j}\|_2 = \sqrt{\|\hat{\boldsymbol{x}}_{\cdot j}\|_2^2 - \sum_{i \in \mathcal{M}(\cdot, j)} (\tilde{x}_{ij}^2 - \hat{x}_{ij}^2)}$, $\|\tilde{\boldsymbol{x}}_{\cdot j}\|_2$ is computed with $\mathcal{O}(|\mathcal{M}(\cdot, j)|)$ time complexity and form the assumptions, the computational cost of $\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)$ is $\mathcal{O}(|\mathcal{M}|)$ and so $L_D(\tilde{w}_j^*), L_P(\tilde{w}_j^*), L_D(\tilde{\alpha}_i^*), L_p(\tilde{\alpha}_i^*)$ are evaluated with $\mathcal{O}(|\mathcal{M}|)$ time complexity.

(ii) The case where the penalty function has **Property B**:

From **Property B**, $\tilde{P}$ is λ-strongly convex and so $\tilde{\boldsymbol{w}}^* \in \Theta_{\tilde{\boldsymbol{w}}^*} := \{\boldsymbol{w} \mid \|\boldsymbol{w} - \hat{\boldsymbol{w}}^*\|_2 \leq \sqrt{2\lambda^{-1}(\tilde{P}(\hat{\boldsymbol{w}}) - \tilde{D}(\hat{\boldsymbol{\alpha}}))}\}$ by using Lemma 2, thereby

$$\hat{w}_j^* - \sqrt{2\lambda^{-1}\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)} \leq \tilde{w}_j^* \leq \hat{w}_j^* + \sqrt{2\lambda^{-1}\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)}.$$

Thus,

$$\begin{aligned}
\tilde{w}_j^* &\geq \hat{w}_j^* - \sqrt{2\lambda^{-1}\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)}, \\
\tilde{w}_j^* &\leq \hat{\alpha}_i^* + \sqrt{2\lambda^{-1}\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)}.
\end{aligned}$$

Similarly, since $\Theta_{\tilde{\boldsymbol{w}}^*}$ is a sphere, the lower and upper bounds of the inner product of arbitrary vector

$\boldsymbol{\eta} \in \mathbb{R}^d$ and $\tilde{\boldsymbol{w}}^*$ are given in closed form as follows by using Lemma 3,

$$\boldsymbol{\eta}^\top \tilde{\boldsymbol{w}}^* \geq \boldsymbol{\eta}^\top \hat{\boldsymbol{w}}^* - \|\boldsymbol{\eta}\|_2 \sqrt{2\lambda^{-1}(\tilde{P}(\hat{\boldsymbol{w}}^*) - \tilde{D}(\hat{\boldsymbol{\alpha}}^*))}, \tag{13}$$

$$\boldsymbol{\eta}^\top \tilde{\boldsymbol{w}}^* \leq \boldsymbol{\eta}^\top \hat{\boldsymbol{w}}^* + \|\boldsymbol{\eta}\|_2 \sqrt{2\lambda^{-1}(\tilde{P}(\hat{\boldsymbol{w}}^*) - \tilde{D}(\hat{\boldsymbol{\alpha}}^*))}. \tag{14}$$

From KKT condition $\tilde{\alpha}_i^* \in -\partial\phi(\tilde{\boldsymbol{z}}_{i\cdot}^\top \tilde{\boldsymbol{w}}^*)$, (10), (13) and since $\partial\phi$ is a monotonically increasing function, $\tilde{\alpha}_i^*$ is bounded as:

$$\begin{aligned}
\tilde{\alpha}_i^* &\geq \inf\left\{-\partial\phi\left(\tilde{\boldsymbol{z}}_{i\cdot}^\top \tilde{\boldsymbol{w}}^*\right)\right\} \\
&\geq \inf\left\{-\partial\phi\left(\max_{\boldsymbol{w}\in\Theta_{\tilde{\boldsymbol{w}}^*}} \tilde{\boldsymbol{z}}_{i\cdot}^\top \boldsymbol{w}\right)\right\} \\
&= \inf\left\{-\partial\phi\left(\tilde{\boldsymbol{z}}_{i\cdot}^\top \hat{\boldsymbol{w}}^* + \|\tilde{\boldsymbol{x}}_{i\cdot}\|_2 \sqrt{2\lambda^{-1}\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)}\right)\right\}
\end{aligned}$$

From KKT condition $\tilde{\alpha}_i^* \in -\partial\phi(\tilde{\boldsymbol{z}}_{i\cdot}^\top \tilde{\boldsymbol{w}}^*)$, (10), (14) and since $\partial\phi$ is a monotonically increasing function, $\tilde{\alpha}_i^*$ is bounded as:

$$\begin{aligned}
\tilde{\alpha}_i^* &\leq \sup\left\{-\partial\phi\left(\tilde{\boldsymbol{z}}_{i\cdot}^\top \tilde{\boldsymbol{w}}^*\right)\right\} \\
&\leq \sup\left\{-\partial\phi\left(\min_{\boldsymbol{w}\in\Theta_{\tilde{\boldsymbol{w}}^*}} \tilde{\boldsymbol{z}}_{i\cdot}^\top \boldsymbol{w}\right)\right\} \\
&= \sup\left\{-\partial\phi\left(\tilde{\boldsymbol{z}}_{i\cdot}^\top \hat{\boldsymbol{w}}^* - \|\tilde{\boldsymbol{x}}_{i\cdot}\|_2 \sqrt{2\lambda^{-1}\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)}\right)\right\}
\end{aligned}$$

The all quantities required for evaluating $L_P(\tilde{w}_j^*), U_P(\tilde{w}_j^*), L_P(\tilde{\alpha}_i^*)$ and $U_P(\tilde{\alpha}_i^*)$ are stored in the memory except $\|\tilde{\boldsymbol{x}}_{i\cdot}\|_2$ and $\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)$. Since $\|\hat{\boldsymbol{x}}_{i\cdot}\|_2$ is stored in the memory and $\|\tilde{\boldsymbol{x}}_{i\cdot}\|_2 = \sqrt{\|\tilde{\boldsymbol{x}}_{i\cdot}\|_2^2 - \sum_{j\in\mathcal{M}(i,\cdot)}(\tilde{x}_{ij}^2 - \hat{x}_{ij}^2)}$, $\|\tilde{\boldsymbol{x}}_{\cdot j}\|_2$ is computed with $\mathcal{O}(|\mathcal{M}(\cdot, j)|)$ and from the assumptions, the computational cost of $\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)$ is $\mathcal{O}(|\mathcal{M}|)$ and so $L_D(\tilde{w}_j^*), L_P(\tilde{w}_j^*), L_D(\tilde{\alpha}_i^*), L_p(\tilde{\alpha}_i^*)$ are evaluated with $\mathcal{O}(|\mathcal{M}|)$ time complexity.

(*iii*) The case where the loss function has **Property A** and the penalty function has **Property B**:

By using **Property A** and **Property B**, $L_D(\tilde{w}_j^*)$, $U_D(\tilde{w}_j^*)$, $L_D(\tilde{\alpha}_i^*)$, $U_D(\tilde{\alpha}_i^*)$, $L_P(\tilde{w}_j^*)$, $U_P(\tilde{w}_j^*)$, $L_P(\tilde{\alpha}_i^*)$ and $U_P(\tilde{\alpha}_i^*)$ can be evaluated with $\mathcal{O}(|\mathcal{M}|)$ time complexity. Thus, $L_{PD}(\tilde{w}_j^*), U_{PD}(\tilde{w}_j^*), L_{PD}(\tilde{\alpha}_i^*), U_{PD}(\tilde{\alpha}_i^*)$ can be evaluated with $\mathcal{O}(|\mathcal{M}|)$ time complexity. □

A.2 Proof of Theorem 2

Proof. Consider partially optimizing the primal problem only w.r.t. a subset of primal solutions $\{\hat{\boldsymbol{w}}_j\}_{j\in\mathcal{J}\subseteq[d]}$. Since

$$\check{\boldsymbol{w}}^*_{\mathcal{J}} = \underset{\{w_j\}_{j\in\mathcal{J}}}{\operatorname{argmin}} \frac{1}{n}\sum_{i\in[n]} \phi\left(\sum_{j\in\mathcal{J}} \tilde{z}_{ij} w_j + \sum_{j\in[d]\setminus\mathcal{J}} \tilde{z}_{ij} w_j\right) + \sum_{j\in\mathcal{J}} \psi_j(w_j),$$

unless $\check{w}^*_j = \hat{w}^*_j$ for all $j \in \mathcal{J}$, because of the strong convexity of $\tilde{P}$, $\tilde{P}(\hat{\boldsymbol{w}}^*) > \tilde{P}(\check{\boldsymbol{w}}^*)$ holds and therefore $\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*) > \tilde{G}(\check{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)$. Thus, both the radius of $\Theta_{\check{\boldsymbol{w}}^*}$ and $\Theta_{\tilde{\boldsymbol{\alpha}}^*}$ (the same definition as in Theorem 1) based on the primal-dual feasible solution pair $(\check{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)$ is smaller than the radius of $\Theta_{\tilde{\boldsymbol{w}}^*}$ and $\Theta_{\tilde{\boldsymbol{\alpha}}^*}$ based on the primal-dual feasible solution pair $(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)$. Similarly, consider partially optimizing the dual problem only w.r.t. a subset of dual solutions $\{\hat{\alpha}_i\}_{i\in\mathcal{I}\subseteq[n]}$. Since

$$\check{\boldsymbol{\alpha}}^*_{\mathcal{I}} = \underset{\{\alpha_i\}_{i\in\mathcal{I}}}{\operatorname{argmax}} -\frac{1}{n}\sum_{i\in\mathcal{I}} \phi^*(-\alpha_i) - \sum_{j\in[d]} \psi^*_j\left(\frac{1}{n}\sum_{i\in\mathcal{I}} \tilde{z}_{ij}\alpha_i + \frac{1}{n}\sum_{i\in[n]\setminus\mathcal{I}} \tilde{z}_{ij}\alpha_i\right),$$

unless $\check{\alpha}^*_i = \hat{\alpha}^*_i$ for all $i \in \mathcal{I}$, because of the strong concavity of $\tilde{D}$, $\tilde{D}(\hat{\boldsymbol{\alpha}}^*) < \tilde{D}(\check{\boldsymbol{\alpha}}^*)$ holds and therefore $\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*) > \tilde{G}(\hat{\boldsymbol{w}}^*, \check{\boldsymbol{\alpha}}^*)$. Thus, both the radius of $\Theta_{\tilde{\boldsymbol{w}}^*}$ and $\Theta_{\tilde{\boldsymbol{\alpha}}^*}$ based on the primal-dual feasible solution pair $(\hat{\boldsymbol{w}}^*, \check{\boldsymbol{\alpha}}^*)$ is smaller than the radius of $\Theta_{\tilde{\boldsymbol{w}}^*}$ and $\Theta_{\tilde{\boldsymbol{\alpha}}^*}$ based on the primal-dual feasible solution pair $(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)$. □

B Bounds in Theorem 1 for smoothed-hinge SVMs

In the case of smoothed-hinge SVM, the bounds in Theorem 1 are written as

$$L_D(\tilde{w}^*_j) = \min\left\{n^{-1}\textstyle\sum_{i|\tilde{z}_{ij}<0} \tilde{z}_{ij}, n^{-1}\lambda^{-1}\tilde{\boldsymbol{z}}_{\cdot j}^\top \hat{\boldsymbol{\alpha}}^* - \lambda^{-1}\|\tilde{\boldsymbol{x}}_j\|_2\sqrt{2\gamma^{-1}\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)}\right\}$$
$$U_D(\tilde{w}^*_j) = \max\left\{n^{-1}\textstyle\sum_{i|\tilde{z}_{ij}>0} \tilde{z}_{ij}, n^{-1}\lambda^{-1}\tilde{\boldsymbol{z}}_{\cdot j}^\top \hat{\boldsymbol{\alpha}}^* + \lambda^{-1}\|\tilde{\boldsymbol{x}}_j\|_2\sqrt{2\gamma^{-1}\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)}\right\}$$

$$L_D(\tilde{\alpha}^*_i) = \min\left\{0, \hat{\alpha}_i - \sqrt{2n\gamma^{-1}\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)}\right\}, \quad U_D(\tilde{\alpha}^*_i) = \max\left\{1, \hat{\alpha}_i + \sqrt{2n\gamma^{-1}\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)}\right\},$$

$$L_P(\tilde{w}^*_j) = \hat{w}_j - \sqrt{2\lambda^{-1}\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)} \quad U_P(\tilde{w}^*_j) = \hat{w}_j + \sqrt{2\lambda^{-1}\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)},$$

$$L_P(\tilde{\alpha}_i^*) = \begin{cases} 0, & (\tilde{\boldsymbol{z}}_{i\cdot}^\top \hat{\boldsymbol{w}}^* + \|\tilde{\boldsymbol{x}}_{i\cdot}\|_2 \sqrt{2\lambda^{-1}\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)} > 1), \\ 1, & (\tilde{\boldsymbol{z}}_{i\cdot}^\top \hat{\boldsymbol{w}}^* + \|\tilde{\boldsymbol{x}}_{i\cdot}\|_2 \sqrt{2\lambda^{-1}\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)} < 1 - \gamma), \\ -\gamma^{-1}(1 - \tilde{\boldsymbol{z}}_{i\cdot}^\top \hat{\boldsymbol{w}}^* - \|\tilde{\boldsymbol{x}}_{i\cdot}\|_2 \sqrt{2\lambda^{-1}\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)}), & (\text{otherwise}), \end{cases}$$

$$U_P(\tilde{\alpha}_i^*) = \begin{cases} 0, & (\tilde{\boldsymbol{z}}_{i\cdot}^\top \hat{\boldsymbol{w}}^* - \|\tilde{\boldsymbol{x}}_{i\cdot}\|_2 \sqrt{2\lambda^{-1}\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)} > 1), \\ 1, & (\tilde{\boldsymbol{z}}_{i\cdot}^\top \hat{\boldsymbol{w}}^* - \|\tilde{\boldsymbol{x}}_{i\cdot}\|_2 \sqrt{2\lambda^{-1}\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)} < 1 - \gamma), \\ -\gamma^{-1}(1 - \tilde{\boldsymbol{z}}_{i\cdot}^\top \hat{\boldsymbol{w}}^* + \|\tilde{\boldsymbol{z}}_{i\cdot}\|_2 \sqrt{2\lambda^{-1}\tilde{G}(\hat{\boldsymbol{w}}^*, \hat{\boldsymbol{\alpha}}^*)}), & (\text{otherwise}). \end{cases}$$